arXiv:1511.08528v1 [math.NA] 26 Nov 2015

GAUSSIAN ELIMINATION WITH RANDOMIZED COMPLETE PIVOTING

CHRISTOPHER MELGAARD* AND MING GU†

Abstract. Gaussian elimination with partial pivoting (GEPP) has long been among the most widely used methods for computing the LU factorization of a given matrix. However, this method is also known to fail for matrices that induce large element growth during the factorization process. In this paper, we propose a new scheme, Gaussian elimination with randomized complete pivoting (GERCP) for the efficient and reliable LU factorization of a given matrix. GERCP satisfies GECP (Gaussian elimination with complete pivoting) style element growth bounds with high probability, yet costs only marginally higher than GEPP. Our numerical experimental results strongly suggest that GERCP is as reliable as GECP and as efficient as GEPP for computing the LU factorization.

1. Background and Motivation.

Solving linear systems of equations

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \tag{1.1}$$

where $\mathbf{A} \in \mathbb{R}^{n\times n}$ and $\mathbf{x}, \mathbf{b} \in \mathbb{R}^n$, is a fundamental problem in numerical linear algebra and scientific computing. Gaussian Elimination with Partial Pivoting (GEPP) solves this problem by computing the LU factorization of $\mathbf{A}$ and is typically efficient and reliable. Over the years, GEPP has been repeatedly re-designed and re-implemented for better performance, and is the backbone for generations of mathematical software packages, including LINPACK [6], LAPACK [1], PLAPACK, SCALAPACK, PLASMA and MAGMA. GEPP routines in today's mathematical software libraries such as the Intel `mkl` are capable of solving linear systems of equations with tens of thousands of variables at or near the peak of the machine's speed.

Efficiency aside, an equally important consideration is numerical reliability. While algorithms for solving eigenvalue problems have become significantly more stable over the years, GEPP was known to be, and remains, a method that is mostly stable in practice but unstable for many well-known matrices including some from common integral equations and differential equations applications [8, 30].

Pivoting plays a crucial role in the reliability of Gaussian elimination (GE), which is tied to *element growth* within the LU factorization process. The most naive version of GE, Gaussian elimination without pivoting (GENP), does not perform any pivoting and only requires $\frac{2}{3}n^3 + O(n^2)$ floating point operations with no entry comparisons [4]. However, this method can suffer from uncontrolled element growth and is only known to be reliable in a few instances like diagonally dominant matrices among others. The most popular version of GE is GEPP, which limits element growth to at most exponential by swapping the rows of $\mathbf{A}$ (i.e., partial pivoting) during elimination, and is numerically stable on average. The additional cost, about $\frac{1}{2}n^2$ entry comparisons and the associated data movement, is typically a small fraction of the total GE cost. The most reliable version of GE is Gaussian elimination with complete pivoting (GECP), which swaps both rows and columns for sub-exponential element growth [28] and is universally believed to be always backward stable in practice [4]. However, GECP is prohibitively slow with $\frac{1}{3}n^3 + O(n^2)$ entry comparisons and relatively little memory reuse [4, 14].

Rook pivoting [9, 21, 24] is an attempt to speedup complete pivoting while maintaining the guarantee of sub-exponential element growth. Rook pivoting is part of the `LUSOL` package [10] for sparse LU factorization. Despite having better performance in the "average" case, there are many matrices that still require $O(n^3)$ entry comparisons in the worst case, providing a negligible speedup over complete pivoting [14].

In this paper, we propose a novel pivoting scheme called *Gaussian elimination with randomized complete pivoting* (GERCP). We show that GERCP satisfies a stability condition similar to that of complete pivoting, suggesting that these methods share similar stability properties. Yet, we also

*Department of Mathematics, University of California, Berkeley.
†Department of Mathematics, University of California, Berkeley. This research was supported in part by NSF Award CCF-1319312.

demonstrate that GERCP is comparable to GEPP in computational overhead. Our numerical experimental results strongly suggest that GERCP is a numerically stable and computationally efficient alternative to GEPP.

Randomization has been used to fix the numerical instability of GEPP in the literature, through GE on the product of random matrices and $\mathbf{A}$ to avoid catastrophically bad pivots. These methods are known to work well in practice in general, but they still lack effective control on element growth, and can be much less accurate than GEPP.

In Section 2, we introduce the necessary notation and background for the paper. In Section 3, we introduce GERCP and state/prove some important properties. In section 4, we talk about numerical experiements and implementations of GERCP. The appendix has results needed by the proofs in section 3.

2. The Setup and Background. In this paper, we consider Gaussian elimination on an invertible square matrix $\mathbf{A} \in \mathbb{R}^{n\times n}$, although our algorithms and analysis carry over to the cases of singular matrices and rectangular matrices with few modifications.

2.1. Notation. We will follow the familiar slight abuse of notation from scientific computing and numerical linear algebra, mimicking the way that LAPACK overwrites the input matrix with the L and U factors. The diagonal of $\mathbf{A}$ becomes the diagonal of U because the diagonal of L is always 1 and thus does not need to be stored.

ALGORITHM 2.1. **Classical Gaussian Elimination in Matlab Notation**

Input: $n \times n$ *matrix* $\mathbf{A}$
Output: *lower triangular* L *with unit diagonal, upper triangular* U*,*
row permutation Π_r*, column permutation* Π_c*.*

---

set $A = \mathbf{A}$
for $k = 1, \cdots, n-1$ **do** *(i.e. called* k^{th} *stage of LU)*

1. **select** *column pivot (INSERT PIVOTING RULE).*
2. **swap** *(UPDATE A AND* Π_c *WITH PIVOT DECISION).*
3. **select** *row pivot (INSERT PIVOTING RULE).*
4. **swap** *(UPDATE A AND* Π_r *WITH PIVOT DECISION).*
5. **compute** $A(k+1:n,k) = A(k+1:n,k)/A(k,k)$*;*
6. **compute** $A(k+1:n,k+1:n) = A(k+1:n,k+1:n) - A(k+1:n,k) * A(k,k+1:n)$*;*

REMARK 2.1. While the working matrix A has been overwritten in Algorithm 2.1, in our subsequent discussions we will refer L and U as the triangular matrices stored in A and still refer $\mathbf{A}$ as the original input matrix. We will use $A_k \in \mathbb{R}^{n\times n}$ to refer explicitly to the working matrix before the k^{th} stage of the outer most loop. Thus, A_n refers to the working matrix after the algorithm terminates, i.e. *a*fter the $(n-1)^{th}$ stage of the outer loop.

REMARK 2.2. For easy of discussion, we have written Algorithm 2.1 in such a way that, for each k, it performs possible column pivoting before any possible row pivoting. GENP, GEPP, GECP and rook pivoting can all be written in this form.

For any appropriate dimension m, we denote $\mathbf{e}_i \in \mathbb{R}^m$ to be the i^{th} standard basis vector, i.e. a vector with all entries equal to 0 except for the i^{th} entry which equals 1; we also denote $\mathbf{e} \in \mathbb{R}^m$ to be the vector with all entries equal to 1. Any permutation matrix $\Pi \in \mathbb{R}^{n\times n}$ is a square matrix with exactly one entry equal to 1 in each row and column, and all other entries equal to 0. We refer to the permutation induced by Π as $\pi : \{1, \cdots, n\} \to \{1, \cdots, n\}$ in the sense that $\pi(i) = j$ if and only if $\Pi\mathbf{e}_i = \mathbf{e}_j$. We will commonly make use of the swap or 2-*cycle* permutation given by $\pi_{(i,j)}$ or $\Pi_{(i,j)}$ in

Common examples of Matlab notation for $1 \le i \le p \le m$ and $1 \le j \le q \le n$			
Notation	**Pivoted Notation**	**Dimensions**	**Description**
$B(:,:)$	$B(\pi_1(:), \pi_2(:))$	$\mathbb{R}^{n \times n}$	Entire matrix B or $\Pi_1 B \Pi_2^T$ resp.
$B(i,:)$	$B(\pi_1(i), \pi_2(:))$	row vector in $\mathbb{R}^n$	i^{th} row of B or $\Pi_1 B \Pi_2^T$ resp.
$B(:,j)$	$B(\pi_1(:), \pi_2(j))$	column vector in $\mathbb{R}^n$	j^{th} column of B or $\Pi_1 B \Pi_2^T$ resp.
$B(i, j:q)$	$B(\pi_1(i), \pi_2(j:q))$	row vector in $\mathbb{R}^{q-j+1}$	j^{th} through q^{th} entries of i^{th} row of B or $\Pi_1 B \Pi_2^T$ resp.
$B(i:p, j)$	$B(\pi_1(i:p), \pi_2(j))$	column vector in $\mathbb{R}^{p-i+1}$	i^{th} through p^{th} entries of j^{th} column of B or $\Pi_1 B \Pi_2^T$ resp.
$B(i:p, j:q)$	$B(\pi_1(i:p), \pi_2(j:q))$	$\mathbb{R}^{(p-i+1)\times(q-j+1)}$	Submatrix from intersection i^{th} through p^{th} rows and j^{th} through q^{th} columns of B or $\Pi_1 B \Pi_2^T$ resp.

FIG. 2.1. *Table of Matlab notations*

matrix form defined by

$$\pi_{(i,j)}(i) = j, \qquad \pi_{(i,j)}(j) = i \qquad \text{and} \qquad \pi_{(i,j)}(k) = k, \text{ for all } k \neq i, j$$

We denote the final row and column permutations of an algorithm as Π_r and Π_c respectively. At the k^{th} stage of LU, Algorithm 2.1 will swap the k^{th} column with the α_k^{th} column and the k^{th} row with the β_k^{th} row. As a result, we can write

$$\Pi_c = \Pi_{(n-1,\alpha_{n-1})} \cdots \Pi_{(2,\alpha_2)} \Pi_{(1,\alpha_1)}, \quad \Pi_r = \Pi_{(n-1,\beta_{n-1})} \cdots \Pi_{(2,\beta_2)} \Pi_{(1,\beta_1)}$$

as a product of the individual column/row swaps. Furthermore, we define the next notation to give us the first $k-1$ swaps and the last $n-k$ swaps

$$\Pi_{c,k} = \Pi_{(k-1,\alpha_{k-1})} \cdots \Pi_{(2,\alpha_2)} \Pi_{(1,\alpha_1)}$$
$$\Pi_{c,-k} = \Pi_{(n-1,\alpha_{n-1})} \cdots \Pi_{(k+1,\alpha_{k+1})} \Pi_{(k,\alpha_k)}$$

Also, we will use the analogous definition for $\Pi_{r,k}$ and $\Pi_{r,-k}$.

In Figure 2.1, we describe the use of the MATLAB colon notation in combination with the permutations above to explain our rows/columns reorderings of a matrix and its select submatrices.

Let π_c and π_r be the permutation of columns and rows performed by LU respectively. We use the following notation to refer to matrices with the final pivoting applied apriori

$$\mathbf{A}^{\Pi_c} = \Pi_r A \Pi_c^T = A(\pi_r(:), \pi_c(:))$$
$$A_k^{\Pi_c} = \Pi_{r,-k} A \Pi_{c,-k}^T = A_k(\pi_{r,-k}(:), \pi_{c,-k}(:))$$

We do this because all of the pivoting methods discussed in this paper are *top-heavy* as defined in Definition 2.10. The row pivots π_r of a top-heavy pivoting strategy are deterministic given the column pivots π_c applied to $\mathbf{A}$ by the LU factorization because each row pivot must satisfy equation (2.6). Therefore, when writing $\mathbf{A}^{\Pi_c}$, it is understood that the row pivots are the unique set of top-heavy row pivots.

Schur complements form a crucial role in Gaussian elimination. We establish notation for Schur complements as $S_k \in \mathbb{R}^{(k:n)\times(k:n)}$. Notice the use of Matlab notation $(k:n) \times (k:n)$ instead of $(n-k+1) \times (n-k+1)$. The Schur complement S_k will act like a normal $(n-k+1) \times (n-k+1)$ matrix for most operations like matrix multiplication, matrix addition and ect. However, when using the Matlab notation in Figure 2.1 to access entries of S_k, we impose the abuse of notation that rows and columns are enumerated from k to n, instead of 1 to $n-k+1$. For example,

Different A, L and U notations			
Notation	**Dimensions**	**Description**	**Algorithm Pivots**
$\mathbf{A}$, A_1, S_1	$\mathbb{R}^{n\times n}$	Unadulterated input matrix	None
$\mathbf{A}^{\Pi_c}$,$A_1^{\Pi_c}$, $S_1^{\Pi_c}$	$\mathbb{R}^{n\times n}$	Pivoted input matrix	All pivots applied apriori
A	$\mathbb{R}^{n\times n}$	Current working matrix. Overwrite triangular L, U factors and Schur complement in place	Pivots applied as determined by algorithm
A_k	$\mathbb{R}^{n\times n}$	Working matrix at k^{th} stage. Overwrite triangular L_k, U_k factors and Schur complement S_k in place	Pivots applied as determined by algorithm
$A_k^{\Pi_c}$	$\mathbb{R}^{n\times n}$	Working matrix at k^{th} stage, i.e. $\Pi_{r,-k}A_k\Pi_{c,-k}^T$	All pivots applied apriori
S_k	$\mathbb{R}^{(k:n)\times(k:n)}$	k^{th} Schur Complement of $\Pi_{r,k}\mathbf{A}\Pi_{c,k}^T$	Pivots applied as determined by algorithm
$S_k^{\Pi_c}$	$\mathbb{R}^{(k:n)\times(k:n)}$	k^{th} Schur Complement of $\Pi_r\mathbf{A}\Pi_c^T$	All pivots applied apriori

FIG. 2.2. *Table of Notations*

- top left entry of S_k is denoted as $S_k(k,k)$, but NOT $S_k(1,1)$.
- submatrix of last two columns of S_k is denoted by $S_k(:, n-1:n)$, but NOT $S_k(:, n-k : n-k+1)$.

This makes our analysis much cleaner and more straightforward because it synchronizes the enumeration of columns/rows between the k^{th} working matrix A_k and the k^{th} Schur complement S_k which is a submatrix for A_k, i.e. $S_k(k:n, k:n) = A_k(k:n, k:n)$. Given this notation for Schur complements, we formally define the working Schur complement at the k^{th} stage S_k and the fully pivoted k^{th} Schur complement $S_k^{\Pi_c}$

$$S_k(k:n, k:n) = A_k(k:n, k:n)$$
$$S_k^{\Pi_c}(k:n, k:n) = A_k^{\Pi_c}(k:n, k:n)$$

where $A_k \in \mathbb{R}^{n\times n}$ is the working matrix at the k^{th} stage. This implies that $S_k^{\Pi_c}(k:n, k:n) = S_k(\Pi_{r,-k}(k:n), \Pi_{c,-k}(k:n))$, i.e. S_k has the pivots only up to the k^{th} step and S_k^{Π} is already pivoted into the final permutation so that no further pivots are required.

The unit lower triangular matrix $L_{k+1} \in \mathbb{R}^{n\times k}$ and the upper triangular matrix $U_{k+1} \in \mathbb{R}^{k\times n}$ relate to the working matrices A_{k+1}

$$U_{k+1} = \begin{pmatrix} A_{k+1}(1,1) & A_{k+1}(1,2) & \cdots & A_{k+1}(1,k) & \cdots & A_{k+1}(1,n) \\ & A_{k+1}(2,2) & \cdots & A_{k+1}(2,k) & \cdots & A_{k+1}(2,n) \\ & & \ddots & \vdots & \cdots & \vdots \\ & & & A_{k+1}(k,k) & \cdots & A_{k+1}(k,n) \end{pmatrix}$$

$$L_{k+1} = \begin{pmatrix} 1 & & & \\ A_{k+1}(2,1) & 1 & & \\ \vdots & \vdots & \ddots & \\ A_{k+1}(k,1) & A_{k+1}(k,2) & \cdots & 1 \\ A_{k+1}(k{+}1,1) & A_{k+1}(k{+}1,2) & \cdots & A_{k+1}(k{+}1,k) \\ \vdots & \vdots & \vdots & \vdots \\ A_{k+1}(n,1) & A_{k+1}(n,2) & \cdots & A_{k+1}(n,k) \end{pmatrix}$$

for $1 < k \leq n$. Remember that $L = L_{n+1} \in \mathbb{R}^{n\times n}$ and $U = U_{n+1} \in \mathbb{R}^{n\times n}$.

2.2. Useful Preliminary Tools.

2.2.1. Tools from Matrix Analysis. In finite dimensions, it is well known that all normed vector spaces are topologically equivalent. The following lemma gives the equivalence between the $\|\cdot\|_2$ and $\|\cdot\|_\infty$ vector norms, which is easily proven from the definitions of the norms. This result can also be seen as a trivial application of the Fritz John Ellipsoid Theorem [17] for convex regions that are symmetric about the origin.

LEMMA 2.1 (Equivalence of ℓ_2 and ℓ_∞ in finite dimensions). *Let $\mathbf{x} \in \mathbb{R}^d$. Then*

$$\frac{1}{\sqrt{d}} \|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_\infty \leq \|\mathbf{x}\|_2$$

DEFINITION 2.2 (Subordinate matrix norms [14, 15]). *Let $1 \leq p, q \leq \infty$ and let $\mathbf{B} \in \mathbb{R}^{m\times n}$. A* **subordinate matrix norm** *is a matrix norm of the form*

$$\|\mathbf{B}\|_{q,p} = \sup_{0\neq\mathbf{x}\in\mathbb{R}^n} \frac{\|\mathbf{Bx}\|_p}{\|\mathbf{x}\|_q} = \max_{\|\mathbf{x}\|_q=1} \|\mathbf{Bx}\|_p$$

Note that for normal matrix operator norms, we have that $\|\mathbf{B}\|_p = \|\mathbf{B}\|_{p,p}$. We make use of a little known subordinate matrix norm by setting $q = 1$ in the above as in exercise 6.11 of [14].

LEMMA 2.3 (Maximum ℓ_p column norm $\|\cdot\|_{1,p}$). *Let $\mathbf{B} \in \mathbb{R}^{m\times n}$. Then, we have*

$$\|\mathbf{B}\|_{1,p} = \max_{1\leq i\leq n} \|\mathbf{Be}_i\|_p = \max_{1\leq i\leq n} \|\mathbf{B}(:,i)\|_p$$

Proof. Let $0 \neq \mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \in \mathbb{R}^n$. Then

$$\|\mathbf{Bx}\|_p = \left\| \sum_{i=1}^n x_i \mathbf{B}(:,i) \right\|_p \leq \sum_{i=1}^n |x_i| \|\mathbf{B}(:,i)\|_p \leq \left(\max_j \|\mathbf{B}(:,j)\|_p \right) \sum_{i=1}^n |x_i| = \left(\max_j \|\mathbf{B}(:,j)\|_p \right) \|\mathbf{x}\|_1$$

Dividing both sides by $\|\mathbf{x}\|_1$ and taking a supremum, we arrive at

$$\|\mathbf{B}\|_{1,p} \overset{def}{=} \sup_{0\neq\mathbf{x}\in\mathbb{R}^n} \frac{\|\mathbf{Bx}\|_p}{\|\mathbf{x}\|_1} \leq \max_j \|\mathbf{B}(:,j)\|_p$$

Also, let $j^* = \arg\max_j \|\mathbf{B}(:,j)\|_p$ and we arrive at our conclusion by observing

$$\max_j \|\mathbf{B}(:,j)\|_p = \frac{\|\mathbf{Be}_{j^*}\|_p}{\|\mathbf{e}_{j^*}\|_1} \leq \sup_{0\neq\mathbf{x}\in\mathbb{R}^n} \frac{\|\mathbf{Bx}\|_p}{\|\mathbf{x}\|_1} \overset{def}{=} \|\mathbf{B}\|_{1,p}$$

□

In particular, we make frequent use of the two following subordinate matrix norms:

$$\|\mathbf{B}\|_{1,\infty} = \max_{1\leq j\leq n} \|\mathbf{Be}_j\|_\infty = \max_{1\leq i,j\leq n} |\mathbf{B}(i,j)| \qquad \textbf{(Maximum entry norm)} \tag{2.1}$$

$$\|\mathbf{B}\|_{1,2} = \max_{1\leq i\leq n} \|\mathbf{Be}_j\|_2 = \max_{1\leq j\leq n} \|\mathbf{B}(:,j)\|_2 \qquad \textbf{(Maximum } \ell_2 \textbf{ column norm)} \tag{2.2}$$

The next lemma, from line (6.19) in chapter 6.3 of [14], will allow us to control each operator norm $\|\cdot\|_p$ for $1 \leq p \leq \infty$ of our residual error via the largest absolute column and row sum of the residual error.

THEOREM 2.4 (Special case of Riesz-Thorin theorem [14]). *Let* $\mathbf{B} \in \mathbb{R}^{m\times n}$ *and let* $1 \leq p \leq \infty$. *Then*

$$\|\mathbf{B}\|_p \leq \|\mathbf{B}\|_1^{\frac{1}{p}} \|\mathbf{B}\|_\infty^{1-\frac{1}{p}}$$

Setting $p = 2$ in the above theorem, we get a bound on the spectral norm of B. The volume of a parallelepiped (parallelotope) formed from the columns of a matrix $\mathbf{B}$ is given by the absolute value of the determinant of B. A rectangle is formed by forcing the parallelepiped to have only right angles between each of its vectors. The next result states that the volume of parallelepiped (parallelotope) is bounded above by that of the corresponding rectangle (hyperrectangle). This result is crucial to deriving element and column growth factors.

THEOREM 2.5 (Hadamard's Inequality [15]). *Let* $\mathbf{B} \in \mathbb{R}^{m\times m}$. *Then, we have that*

$$|\det(\mathbf{B})| \leq \prod_{j=1}^{m} \|\mathbf{B}(:,j)\|_2$$

2.2.2. Tools from Probability Theory. The union bound is a basic yet important result whose proof is typically left as an exercise in most introductory probability textbooks. We will make use of it in combination with the famous De Morgan laws to consider the probability of an intersection of highly coupled events in the analysis of GERCP.

LEMMA 2.6 (Union Bound or Boole's Inequality). *For events* $E_1, E_2, ..., E_m$, *we have that*

$$\mathbb{P}\left(\bigcup_{i=1}^{m} E_i\right) \leq \sum_{i=1}^{m} \mathbb{P}\left(E_i\right)$$

As GERCP is a randomized algorithm, the factorization it produces will be random. The Law of Total Probability below is the basis on which we analyze the reliability of GERCP regardless of the column and row permutations chosen by GERCP.

THEOREM 2.7 (Law of Total Probability). *Given* m *mutually exclusive events* $E_1, \cdots, E_m$ *whose probabilities sum to unity, then*

$$\mathbb{P}\left(B\right) = \sum_{i=1}^{m} \mathbb{P}\left(B|E_i\right)\mathbb{P}\left(E_i\right),$$

where B *is an arbitrary event, and* $\mathbb{P}\left(B|E_i\right)$ *is the conditional probability of* B *assuming* E_i.

We will also need the following large deviation result, due to Johnson and Lindenstrauss, from probability theory and theoretical computer science. Theorem 2.8 has been the main theoretical foundation in the recent development of randomized algorithms in numerical linear algebra and data analysis. It will also be the primary theoretical justification for the reliability of GERCP.

THEOREM 2.8 (Random Projection Method (Johnson-Lindenstrauss) [26]). *Let* $\mathbf{x} \in \mathbb{R}^d$ *and* $\epsilon > 0$. *Assume that the entries in* $\Omega \in \mathbb{R}^{r\times d}$ *are sampled independently from* $\mathcal{N}(0,1)$. *Then*

$$\mathbb{P}\left((1-\epsilon)\|\mathbf{x}\|_2^2 \geq \left\|\frac{1}{\sqrt{r}}\Omega\mathbf{x}\right\|_2^2\right) \leq \exp\left(-\frac{(\epsilon^2-\epsilon^3)r}{4}\right)$$

$$\mathbb{P}\left(\left\|\frac{1}{\sqrt{r}}\Omega\mathbf{x}\right\|_2^2 \geq (1+\epsilon)\|\mathbf{x}\|_2^2\right) \leq \exp\left(-\frac{(\epsilon^2-\epsilon^3)r}{4}\right)$$

and

$$\mathbb{P}\left((1-\epsilon)\left\|\mathbf{x}\right\|_2^2 \leq \left\|\frac{1}{\sqrt{r}}\Omega\mathbf{x}\right\|_2^2 \leq (1+\epsilon)\left\|\mathbf{x}\right\|_2^2\right) \geq 1-2\exp\left(-\frac{(\epsilon^2-\epsilon^3)r}{4}\right). \tag{2.3}$$

Due to the central importance of Theorem 2.8 in GERCP, we make the next definition

DEFINITION 2.9. *A given vector $\mathbf{x} \in \mathbb{R}^d$ satisfies the $\epsilon-$JL condition under random mapping Ω if*

$$\sqrt{1-\epsilon}\left\|\mathbf{x}\right\|_2 \leq \left\|\frac{1}{\sqrt{r}}\Omega\mathbf{x}\right\|_2 \leq \sqrt{1+\epsilon}\left\|\mathbf{x}\right\|_2.$$

REMARK 2.3. Despite its simplicity, Theorem 2.8 asserts the surprisingly strong norm-preserving abilities under a random projection. For any given $\Delta \in (0,1)$, x satisfies the $\epsilon-$JL condition under random mapping Ω with probability at least $1-\Delta$ for any

$$r \geq \frac{4}{\epsilon^2-\epsilon^3}\log\left(\frac{2}{\Delta}\right). \tag{2.4}$$

In particular, for $\epsilon = \frac{1}{2}$ and $\Delta = 10^{-5}$, $r = 400$ satisfies equation (2.4), regardless of d. In practice, however, one can typically choose a much smaller value of r for x to satisfy $\epsilon-$JL condition.

2.3. Numerical Error and Stability of LU Factorization. With computer roundoff, Theorem 9.3 of [14] gives us that our computed LU factorization obeys the relationship $\mathbf{A}+E = LU$ such that

$$|E_{jk}| \leq \frac{n\epsilon_{mach}}{1-n\epsilon_{mach}}\left(|L|\cdot|U|\right)_{jk},$$

where $E \in \mathbb{R}^{n\times n}$ and $|\cdot|$ denotes taking the absolute value of each entry of a matrix. For linear systems $LUx = b$, the backwards stability of forward/backward substitution tells us that our computer will calculate $\hat{x}$ which satisfies the following approximation

$$(L+\delta L)(U+\delta U)\hat{x} = b$$

Thus, Gaussian elimination for the linear system is backwards stable if we can provide a tight bound on δA such that $(A+\delta A)\hat{x} = b$ where

$$\delta\mathbf{A} = \delta LU + L\delta U + \delta L\delta U + E$$

Theorem 9.4 of [14] tells us that δA must satisfy

$$|\delta\mathbf{A}| \leq \frac{3n\epsilon_{mach}}{1-3n\epsilon_{mach}}|L|\,|U| \tag{2.5}$$

Therefore, in order to bound δA, we need to simply bound L and U of our computed factorization. Next, we define a property that some pivoting strategies enjoy.

DEFINITION 2.10 (Top-Heavy Pivoting Strategies). *A pivoting strategy for Gaussian elimination or the LU decomposition is called* **top-heavy** *if the pivoting strategy leaves the first entry of the leading column of each Schur complement to be the entry with largest modulus in the leading column. In other words,*

$$\left|S_k^{\Pi_c}(k,k)\right| = \max_{k\leq i\leq n}\left|S_k^{\Pi_c}(i,k)\right| \tag{2.6}$$

or, in other words

$$\left|S_k^{\Pi_c}(k,k)\right| = \|S_k^{\Pi_c}(:,k)\|_\infty$$

for all $1 \leq k \leq n$. All of the methods discussed in this paper (partial, complete, rook, ℓ_2 complete and randomized complete pivoting) are **top-heavy** pivoting strategies. Therefore, all of these strategies enjoy the following property

LEMMA 2.11. *Let* $\mathbf{A} \in \mathbb{R}^{n\times n}$ *and let the lower triangular matrix* L *be obtained by the* LU *algorithm above with a top-heavy pivoting strategy. Then, we have*

$$\|L\|_p \leq n$$

where $1 \leq p \leq \infty$.

Proof. First, note that top-heaviness gives us that

$$|l_{jk}| = \frac{|S_k(j,k)|}{|S_k(k,k)|} \leq 1$$

for $k \leq j$. Thus,

$$\|L\|_1 \leq n, \qquad \|L\|_\infty \leq n \quad \text{and} \quad \|L\|_p \leq \|L\|_1^{\frac{1}{p}} \|L\|_\infty^{1-\frac{1}{p}} \leq n$$

where the last inequality was established by Theorem 2.4. ∎

It is easy to see that all top-heavy strategies enjoy the bound $\|L\|_p \leq n$ for $1 \leq p \leq \infty$ because all the entries below the diagonal in L will be between -1 and 1 in addition to Lemma 2.4. Bounding U is a little trickier and historically, the element growth factor has been used to do it.

DEFINITION 2.12 (Element Growth Factor). *Let* $n > 0$ *and* $\mathbf{A} \in \mathbb{R}^{n\times n}$

$$\rho_{elem}(\mathbf{A}) \stackrel{def}{=} \frac{\max_k \|S_k\|_{1,\infty}}{\|\mathbf{A}\|_{1,\infty}} = \frac{\max_{i,j,k} |S_k(i,j)|}{\max_{i,j} |\mathbf{A}(i,j)|}$$

In addition to the classical element growth factor, we define the new column growth factor which will be central to our analysis.

DEFINITION 2.13 (Column Growth Factor). *Let* $n > 0$ *and* $\mathbf{A} \in \mathbb{R}^{n\times n}$

$$\rho_{col}(\mathbf{A}) \stackrel{def}{=} \frac{\max_k \|S_k\|_{1,2}}{\|\mathbf{A}\|_{1,2}} = \frac{\max_{j,k} \|S_k(:,j)\|_2}{\max_j \|\mathbf{A}(:,j)\|_2}$$

These two definitions of the growth factor are related by the following lemma. It is important to note that the column growth factor commonly attains the lower bound of $\frac{1}{\sqrt{n}}\rho_{elem}$ as we will see with partial and complete pivoting; making ρ_{col} a more informative quantity to control than ρ_{elem} by a factor of $\sqrt{n}$.

LEMMA 2.14. *Let* $n > 0$ *and* $\mathbf{A} \in \mathbb{R}^{n\times n}$, *then*

$$\frac{1}{\sqrt{n}}\rho_{elem}(\mathbf{A}) \leq \rho_{col}(\mathbf{A}) \leq \sqrt{n}\rho_{elem}(\mathbf{A})$$

Proof. Easy consequence of Lemma 2.1. ∎

Using the definition of element growth, we can bound U in the $\|\cdot\|_1$ and $\|\cdot\|_\infty$ norms as

$$\|U\|_\eta \leq n\rho_{elem}(\mathbf{A})\|\mathbf{A}\|_{1,\infty} \qquad \text{for } \eta = 1, \infty$$

Plugging both of these estimates into Theorem 2.4, we get that

$$\|U\|_p \leq n\rho_{elem}(\mathbf{A})\|\mathbf{A}\|_{1,\infty} \qquad \text{for } 1 \leq p \leq \infty$$

This gives us the following classical result

THEOREM 2.15 (Wilkinson [4, 14]). *Let $\mathbf{A} \in \mathbb{R}^{n\times n}$ and let E and $\delta\mathbf{A}$ be given above, then for any top-heavy pivoting strategy, we have*

$$\|\delta\mathbf{A}\|_p \leq \frac{3n\epsilon_{mach}}{1-3n\epsilon_{mach}} n^2 \rho_{elem}(\mathbf{A})\|\mathbf{A}\|_{1,\infty}$$
$$\|E\|_\eta \leq \frac{n\epsilon_{mach}}{1-n\epsilon_{mach}} n^2 \rho_{elem}(\mathbf{A})\|\mathbf{A}\|_{1,\infty}$$

where $1 \leq p \leq \infty$. The proof of the above is similar to the proof of the corresponding result for the column growth factor.

THEOREM 2.16 (Column growth control on backward error). *Let $\mathbf{A} \in \mathbb{R}^{n\times n}$ and let E and $\delta\mathbf{A}$ be given above, then for any top-heavy pivoting strategy, we have*

$$\|\delta\mathbf{A}\|_p \leq \frac{3n\epsilon_{mach}}{1-3n\epsilon_{mach}} n^2 \rho_{col}(\mathbf{A})\|\mathbf{A}\|_{1,2}$$
$$\|E\|_p \leq \frac{n\epsilon_{mach}}{1-n\epsilon_{mach}} n^2 \rho_{col}(\mathbf{A})\|\mathbf{A}\|_{1,2}$$

where $1 \leq p \leq \infty$.

Proof. First, Lemma 2.14 gives use the desired bound on $\|L\|_p$. Next, please note that by the definition of the LU algorithm, we have $U(m, m:n) = S_m(\beta_m, \pi_{(m,\alpha_m)}(m:n))$. Next, we tackle the U matrix will ρ_{col}.

$$\begin{aligned}\|U\|_\infty &= \max_{1\leq m\leq n}\|U(m,:)\|_1 \leq \sqrt{n}\max_{1\leq m\leq n}\|U(m,:)\|_2 = \sqrt{n}\max_{1\leq m\leq n}\|S_m(\beta_m,:)\|_2 \\ &= \sqrt{n}\max_{1\leq m\leq n}\|S_m\|_F \leq n\max_{1\leq m\leq n}\max_{m\leq j\leq n}\|S_m(:,j)\|_2 = n\max_{1\leq m\leq n}\|S_m\|_{1,2}\end{aligned}$$

and

$$\begin{aligned}\|U\|_1 &= \max_{1\leq j\leq n}\|U(:,j)\|_1 \leq n\max_{1\leq j\leq n}\|U(:,j)\|_\infty = n\max_{1\leq j\leq n}\max_{1\leq m\leq j}|U(m,j)| \\ &= n\max_{1\leq m\leq n}\max_{m\leq j\leq n}|S_m(m,j)| \leq n\max_{1\leq m\leq n}\max_{m\leq j\leq n}\|S_m(:,j)\|_2 = n\max_{1\leq m\leq n}\|S_m\|_{1,2}\end{aligned}$$

Thus, Theorem 2.4 gives us

$$\|U\|_p \leq \|U\|_1^{\frac{1}{p}}\|U\|_\infty^{1-\frac{1}{p}} \leq n\max_{1\leq m\leq n}\|S_m\|_{1,2}$$

for all $1 \leq p \leq \infty$. Note that

$$\max_{1\leq m\leq n}\|S_m\|_{1,2} = \frac{\max_{1\leq m\leq n}\|S_m\|_{1,2}}{\|\mathbf{A}\|_{1,2}}\|\mathbf{A}\|_{1,2} = \rho_{col}(\mathbf{A})\|\mathbf{A}\|_{1,2}$$

We arrive at our conclusion by combining this with equation (2.5). ∎

2.4. Popular Pivoting Strategies.

2.4.1. No Pivoting or Static Pivoting. The easiest and most computationally efficient strategy is not to pivot at all, which we call *Gaussian elimination with no pivoting* (GENP). However, this method is not numerically stable for a general $\mathbf{A} \in \mathbb{R}^{n\times n}$. To demonstrate this, we use an example from [14]

$$\mathbf{A} = \begin{bmatrix} \delta & -1 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \delta^{-1} & 1 \end{bmatrix} \begin{bmatrix} \delta & -1 \\ 0 & 1+\delta^{-1} \end{bmatrix} = LU$$

where $\delta < \epsilon_{mach}$ and ϵ_{mach} is machine epsilon. However, in floating point arithmetic, we have $fl\left(1+\delta^{-1}\right) = \delta^{-1}$ where $fl(\cdot)$ represents evaluation in floating point arithmetic. Thus,

$$fl\left(LU\right) = fl\left(L\right) fl\left(U\right) = \begin{bmatrix} 1 & 0 \\ \delta^{-1} & 1 \end{bmatrix} \begin{bmatrix} \delta & -1 \\ 0 & \delta^{-1} \end{bmatrix} = \begin{bmatrix} \delta & -1 \\ 1 & 0 \end{bmatrix} \neq \mathbf{A}$$

which is the wrong computed LU factorization with backward error $\|\mathbf{A} - fl(L)fl(U)\|_\infty = 1$. On the other hand, if $\mathbf{A} \in \mathbb{R}^{n\times n}$ is

- Totally nonnegative, i.e. the determinant of every square submatrix is nonnegative,
- Row or column diagonally dominant,
- Symmetric positive definite,

then it is proved in [14] that *no pivoting* is required for a stable computation. In fact, all of these matrices have element growth $\rho_n = O(1)$.

2.4.2. Partial Pivoting. The most common version of LU is *Gaussian elimination with partial pivoting* (GEPP) because it provides some stability at relatively cheap overhead. It is backward stable "in practice" [4], meaning that this method provides a stable LU factorization for most but not all matrices. For partial pivoting, please place the following pivoting rule in Algorithm 1: At the k^{th} stage, the k^{th} row is swapped with the β_k^{th} row, where

$$\beta_k = \underset{k\leq i\leq n}{\arg\max} \left|S_k(i,k)\right|$$

This method requires a number of entry comparisons in addition to the floating point operations required by Gaussian elimination without pivoting. Specifically, it requires

$$\sum_{k=1}^{n}(n-k) = \frac{n(n-1)}{2}$$

comparisons in total, which is one order less asymptotically than the $\frac{1}{3}n^3 + O(n^2)$ flops in GENP. Each row swap requires n individual entry swaps, so the total number of entry swaps required is bounded above by

$$\sum_{k=1}^{n} n = n^2$$

Since data movement is typically far more expensive than comparisons per operation, the dominant overhead cost is typically the data movement involved with swapping. The element growth for partial pivoting is bounded by

$$\rho_{elem}^{gepp}\left(\mathbf{A}\right) \leq 2^{n-1}$$

It is also easy to show that the column growth for partial pivoting is bounded by

$$\rho_{col}^{gepp}\left(\mathbf{A}\right) \leq \frac{1}{\sqrt{n}} 2^{n-1}$$

Both of these bounds are attained by the Wilkinson matrix

$$\begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 1 \\ -1 & 1 & 0 & \cdots & 0 & 1 \\ -1 & -1 & 1 & \cdots & 0 & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ -1 & -1 & -1 & \cdots & 1 & 1 \\ -1 & -1 & -1 & \cdots & -1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ -1 & 1 & 0 & \cdots & 0 & 0 \\ -1 & -1 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ -1 & -1 & -1 & \cdots & 1 & 0 \\ -1 & -1 & -1 & \cdots & -1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 1 \\ 0 & 1 & 0 & \cdots & 0 & 2 \\ 0 & 0 & 1 & \cdots & 0 & 2^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 2^{n-2} \\ 0 & 0 & 0 & \cdots & 0 & 2^{n-1} \end{bmatrix} = LU$$

where we see that the element growth in U is 2^{n-1} and the column growth $\frac{\|2^{n-1}\|_{\ell_2(\mathbb{R}^1)}}{\|\mathbf{e}\|_{\ell_2(\mathbb{R}^n)}} = \frac{2^{n-1}}{\sqrt{n}}$.

Large element growth and unstable LU factorizations with partial pivoting also occur in many applications. Wright [30] describes a family of two-point boundary value problems that cause GEPP to fail via exponential element growth when attempting to solve the ODE by discretizing it into matrix form. Liu and Russell [19] experience the same phenomenon when attempting to solve the discretized Kuramoto-Sivashinsky PDE, which is used to model laminar flame front propagation, phase dynamics in reaction-diffusion systems, fluctuations in fluid films and instabilities in plasma physics [16, 18]. Foster [8] applies the Newton-Cotes quadrature to discretize the Volterra Integral equation from many areas of applied mathematics including actuarial science, viscoelastic materials and probability theory. This reduces the Volterra Integral equation into a matrix equation that makes GEPP fail. As we will discuss later, the Volterra example is among the most diabolical examples that break GEPP because it induces *passive aggressive element growth*, i.e. barely enough element growth to cause GEPP to fail.

The remainder of the partial pivoting section is spent discussing the *generalized Wilkinson Matrix*, which is a more general class of matrices that can cause exponential element growth in GEPP.

EXAMPLE 2.1. [Generalized Wilkinson Matrices $\mathcal{GW}$] For any integer $r \geq 1$, consider a matrix $\mathbf{A}$ of the following form

$$\mathbf{A} = L + \begin{pmatrix} 1 \\ \vdots \\ 1 \\ 0 \end{pmatrix} \begin{pmatrix} 0 & \dots & 0 & 1 \end{pmatrix},$$

where L is a lower triangular matrix with

$$L_{i,i} = 1, \quad \text{and} \quad L_{i,j} = -u_i^T W_{i+1} \cdots W_{j-1} v_j, \quad \text{for any} \quad i > j,$$

with $u_i, v_j \in \mathbb{R}^r$ being any vectors and $W_i \in \mathbb{R}^{r \times r}$ being square matrices. The matrix $\mathbf{A}$ reduces to the Wilkinson matrix for the special case $r = 1$, $u_i = v_j = W_i = 1$ for all i and j. It is straightforward to verify that L^{-1} is a lower triangular matrix with

$$\left(L^{-1}\right)_{i,i} = 1, \quad \text{and} \quad \left(L^{-1}\right)_{i,j} = u_i^T \widehat{W}_{i+1} \cdots \widehat{W}_{j-1} v_j, \quad \text{for any} \quad i > j,$$

where $\widehat{W}_i = W_i + v_i u_i^T$. Now we choose the vectors $\{u_i\}_{i=2}^n$, $\{v_j\}_{j=1}^{n-1}$ and matrices $\{W_i\}_{i=2}^{n-1}$ to contain only positive entries and have 2-norm at most 1. This implies that

$$|L_{i,j}| = |u_i^T W_{i+1} \cdots W_{j-1} v_j| \leq 1 \quad \text{for any} \quad i > j.$$

Consequently LU-factorizing $\mathbf{A}$ with GEPP will incur no row exchanges, and the resulting matrix factorization has the form

$$\mathbf{A} = L\,U, \quad \text{where} \quad U = I + \left(L^{-1} \begin{pmatrix} 1 \\ \vdots \\ 1 \\ 0 \end{pmatrix} \right) \begin{pmatrix} 0 & \cdots & 0 & 1 \end{pmatrix}.$$

This typically implies exponential element growth in U if the inequality $\|\widehat{W}_i\|_2 > 1$ holds for most matrices $\widehat{W}_i$.

In our numerical experiments, we use this to create a random matrix ensemble that causes GEPP to fail with high probability.

2.4.3. Complete Pivoting. The most reliable version is *Gaussian elimination with complete pivoting* (GECP). Von Neumann and Goldstine [27] referred to this as the "customary procedure," making it arguably the first published pivoting strategy for Gaussian elimination in large matrix computations. For complete pivoting, please place the following pivoting rule in Algorithm 1: At k^{th} stage, k^{th} row and k^{th} column are swapped with β_k^{th} row and α_k^{th} column respectively , where

$$(\beta_k, \alpha_k) = \underset{\substack{(i,j)\in\mathbb{N}^2 \\ k\le i,j\le n}}{\arg\max} |S_k(i,j)|$$

Next, we consider the overhead of complete pivoting, which is broken into entry comparisons and data movement. The number of overall entry comparisons is

$$\sum_{k=1}^{n} \left(k^2 - 1\right) = \frac{n(n+1)(2n+1)}{6} - n = \frac{1}{3}n^3 + \frac{1}{2}n^2 - \frac{5}{6}n$$

In the worst case, the number of entry swaps (or data movement) is

$$\sum_{k=1}^{n} 2n = 2n^2$$

Therefore, despite the fact that each entry swap is more expensive on modern computers than each entry comparison, the entry comparisons will form the bulk of the overhead when n is large. This is different than partial pivoting because the total number of comparisons in complete pivoting is $O(n^3)$ instead of the $O(n^2)$ comparisons in partial pivoting. In his seminal work, Wilkinson [28] proves that the element growth in complete pivoting is bounded above by

$$\rho_{elem}^{gecp}(\mathbf{A}) \le \sqrt{n}\left(2 \cdot 3^{\frac{1}{2}} \cdots n^{\frac{1}{n-1}}\right)^{1/2} \sim c n^{1/2} n^{\frac{1}{4}\ln(n)}$$

Our proof of Theorem 3.3 can be easily adapted to show that the column growth for complete pivoting is bounded above by

$$\rho_{col}^{gecp}(A) \le \left(2 \cdot 3^{\frac{1}{2}} \cdots n^{\frac{1}{n-1}}\right)^{1/2} \sim c n^{\frac{1}{4}\ln(n)}$$

These bounds are provably unattainable for $n \ge 3$ because of their proof's reliance on Hadamard's inequality of Theorem 2.5. It was incorrectly conjectured that $\rho_{elem}^{gecp}(\mathbf{A}) \le n$ [7, 11]. Nonetheless, it is widely believed that the above element growth bound is wildly pessimistic. However, this bound proves that exponential element growth is impossible as $n^{\log(n)}$ is sub-exponential. Because of this, we call GECP *backwards stable*.

2.4.4. Rook Pivoting. An important attempt to speed up complete pivoting was introduced by Neal and Poole [21] as *Gaussian elimination with rook pivoting* (GERP). Basically, one alternates between partial pivoting on the rows and the columns until arriving at a Nash-equilibrium of sorts. For rook pivoting, place the following pivoting rule in Algorithm 1: At k^{th} stage, initialize $\beta_k = k$ and $\alpha_k = k$. First, choose a new β_k from (2.7) while holding α_k constant, and then choose α_k from

(2.8) while holding β_k constant. Repeat (2.7) and (2.8) until the current choice (β_k, α_k) make (2.7) and (2.8) hold simultaneously.

$$\beta_k = \underset{k \le i \le n}{\arg\max} |S_k(i, \alpha_k)| \tag{2.7}$$

$$\alpha_k = \underset{k \le j \le n}{\arg\max} |S_k(\beta_k, j)| \tag{2.8}$$

The number of entry comparisons required by rook pivoting depends on the matrix. For example, a diagonally dominant matrix will require no swaps, which causes rook pivoting to only check (2.7) once. This example gives $\sum_{k=1}^{n}(n-k) = \frac{n(n-1)}{2} = O(n^2)$ entry comparisons just like partial pivoting. In fact, there are a few probabilistic arguments [9, 24] that claim for the "average" input matrix $\mathbf{A}$, the user would expect to perform $O(n^2)$ entry comparisons in total. However, the worst case needs a total of $O(n^3)$ entry comparisons (just as bad as complete pivoting) as exemplified by any matrix of the form [14]

$$\begin{pmatrix} \theta_1 & \theta_2 & & & \\ & \theta_3 & \theta_4 & & \\ & & \ddots & \ddots & \\ & & & \theta_{2n-3} & \theta_{2n-2} \\ & & & & \theta_{2n-1} \end{pmatrix}, \qquad |\theta_1| < |\theta_2| < \cdots < |\theta_{2n-1}|$$

The worst case data movement in terms of entry-wise swaps is the same as GECP at $O(n^2)$ because both rows and columns are also being swapped here. It is not clear whether comparisons or swaps are to be considered the dominant overhead cost as it will depend on whether the input matrix requires $O(n^3)$ or $O(n^2)$ comparisons. Foster [9] proves that rook pivoting element growth must obey

$$\rho_{elem}^{gerp}(\mathbf{A}) \le \frac{3}{2} n^{\frac{3}{4}\ln(n)}$$

and he also shows that the bound is unattainable for $n \ge 3$. This suggests that this method has similar stability properties to complete pivoting and should be considered as a less expensive "cousin."

2.4.5. Prior attempts at randomizing Gaussian elimination. Randomization in the context of Gaussian elimination based direct solvers has been attempted in the past as a way to avoid pivoting altogether [22, 23]. These methods pre/post multiply our input matrix $\mathbf{A}$ by random matrices before applying Gaussian elimination without pivoting. This can either be used to solve a linear system or compute the inverse of $\mathbf{A}$. Unfortunately, there are no theoretical guarantees of small backwards error. In fact, many examples, including the Turing matrix, cause these methods to out-right fail in terms of backwards error as shown in the numerical experiments section of [23].

3. Algorithm GERCP and Main Results. In this section, we first introduce Algorithm 3.1, a deterministic complete pivoting scheme based on the ℓ_2 norm; we then evolve Algorithm 3.1 into GERCP by significantly speeding it up with randomization.

3.1. A Deterministic ℓ_2-norm Complete Pivoting Strategy. To help motivate our randomized strategy, we propose an intermediate deterministic strategy called *Gaussian elimination with ℓ_2 complete pivoting* (GE2CP).

ALGORITHM 3.1. **Gaussian Elimination with ℓ_2 norm Complete Pivoting (GE2CP)**

Input:	*$n \times n$ matrix* $\mathbf{A}$
Output:	*lower triangular L, upper triangular U, row permutation Π_r, column permutation Π_c.*

for $k = 1, \cdots, n-1$ **do**

1. **compute** $\alpha = \mathbf{argmax}_{k \le j \le n} \|A(k:n, j)\|_2$.
2. **swap** *columns* k *and* α *of* A.
3. **compute** $\beta = \mathbf{argmax}_{k \le j \le n} |A(j, k)|$.
4. **swap** *rows* k *and* β *of* A.
5. **compute** $A(k+1:n, k) = A(k+1:n, k)/A(k, k)$;
6. **compute** $A(k+1:n, k+1:n) = A(k+1:n, k+1:n) - A(k+1:n, k) * A(k, k+1:n)$;

Pivoting is done in two steps in Algorithm 3.1 for each k: Step 1 swaps the ℓ^{th} column of the trailing matrix $A(k:n, k:n)$ with the k^{th}, where $A(k:n, \alpha)$ has the largest column 2-norm among all columns of $A(k:n, k:n)$; whereas Step 2 swaps the β^{th} row of $A(k:n, k:n)$ with the k^{th}, where $A(\beta, k)$ is the largest in absolute value among all entries of $A(k:n, k)$. Step 1 controls potentially harmful column norm growth in $\mathbf{A}$ through column interchanges, and Step 2 performs standard partial pivoting to control potentially harmful element growth in the k-column $A(k:n, k)$ through row interchanges. Step 2 makes the ℓ_2 norm Complete Pivoting strategy (GE2CP) in Algorithm 3.1 a *top-heavy* pivoting strategy, which allows us to apply Theorems 2.16 to bound the LU backward error.

In this method, the number of comparisons is reduced in favor of additional floating point operations. The total amount of comparisons for this method is

$$\sum_{k=1}^{n} 2(n-k) = n(n-1) = O(n^2).$$

The total additional floating point operations required to directly compute the 2-norms in Step 1 of Algorithm 3.1 is about

$$\sum_{k=1}^{n-1} 2(n-k+1)(n-k+1) = \frac{n(n+1)(2n+1)}{3} - 2 = O(n^3).$$

The worst case for entry swaps in GE2CP is the same as in GECP at $O(n^2)$. The ℓ_2 norm complete pivoting strategy (GE2CP) in Algorithm 3.1 enjoys similar element/column growth bounds to complete pivoting and rook pivoting. By omitting references and applications of Lemma 3.2 from the proof of Theorem 3.3, one can easily show that

$$\rho_{col}^{ge2cp} \le n\sqrt{e(n+1)}n^{\frac{1}{2}\ln(n)}$$

The $O(n^3)$ flop overhead makes Algorithm 3.1 an impractical alternative to GEPP. In Section 3.2, we develop GERCP by randomizing Algorithm 3.1 to choose columns with sufficiently large column norms at significantly lowered overhead costs. Furthermore, we will derive a GECP style element growth upper bound for GERCP that holds with arbitrarily high probability.

3.2. Gaussian Elimination with Randomized Complete Pivoting (GERCP). Our randomized Gaussian Elimination algorithm, GERCP, is based, in principle, on Algorithm 3.1. However, the key difference is we will replace the column pivoting step, Step 1, by a randomized alternative to significantly reduce its cost. At its core, GERCP relies on a fast random project scheme to reliably estimate the column norms of *each* observed Schur complement required by line 1 of Algorithm 3.1, based on Theorem 2.8,

3.2.1. Successive Schur Sketching. We adapt the idea of a *sketching matrix* [29] in this section to speed up the column selection procedure in GE2CP Algorithm 3.1. Let our sampling matrix $\Omega \in \mathbb{R}^{r \times n}$ be a random matrix with iid standard normal $\mathcal{N}(0, 1)$ entries where $r \ll n$, and let

$\Psi = \Omega A$ for our first sketching matrix. We refer to the number $r \in \mathbb{N}$ as the **sampling dimension**. One can also make use of a Fast Johnson-Lindenstrauss-like sampling matrix **??** as a sampling matrix. However, we will stick with a Gaussian sampling matrix in this paper for simplicity of presentation. By Theorem 2.8, the column norms of A can be reliably estimated via those of Ψ for a large enough choice of r. In other words, we can perform Step 1 of Algorithm 3.1 for $k = 1$ by looking for the column with the largest column norm in Ψ.

However, Step 1 of Algorithm 3.1 must be performed for every value of k. From the single random matrix Ω, below we will construct a collection of matrices $\{\Psi_k\}_{k=1}^{n}$ known as the **Schur sketching matrices**, where $\Psi_k \in \mathbb{R}^{r\times(n-k+1)}$. These matrices are constructed as

$$\Psi_k \stackrel{def}{=} \Omega\left(:, \pi_{r,k}(k:n)\right)\ S_k, \text{ for all } 1 < k \leq n. \tag{3.1}$$

REMARK 3.1. *For randomized complete pivoting, the choice of the k^{th} pivot column will be based on column norms in the Schur sketching matrix Ψ_k. Since all Schur sketching matrices are based on the same random matrix Ω, the Schur complement $S_k \in \mathbb{R}^{(k:n)\times(k:n)}$ for $1 < k \leq n$ will* **not** *be a deterministic matrix. Indeed, the observed Schur complement S_k is determined by the randomized column pivoting decisions $\alpha_1, \cdots, \alpha_{k-1}$ of GERCP from the previous stages. Given that there are only a finite number of pivot decisions, we conclude that S_k must be a discrete random variable.*

To efficiently continue with all other column pivoting work in Algorithm 3.1, we inductively devise a scheme to use our current Schur sketching matrix Ψ to produce our next Schur sketching matrix $\widehat{\Psi}$. Suppose that we have chosen $p > 0$ rows/columns from the remaining m rows/columns to pivot to the front of our working matrix

$$\begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} = \begin{pmatrix} L_{11} & \\ L_{21} & I \end{pmatrix} \begin{pmatrix} U_{11} & U_{12} \\ & \widehat{A}_{22} \end{pmatrix},$$

where $A_{11}, L_{11}, U_{11} \in \mathbb{R}^{p\times p}$, with L_{11} lower triangular and U_{11} upper triangular, respectively; $A_{21}, L_{21} \in \mathbb{R}^{(m-p)\times p}$; $A_{12}, U_{12} \in \mathbb{R}^{p\times(m-p)}$; and $A_{22}, \widehat{A}_{22} \in \mathbb{R}^{(m-p)\times(m-p)}$ with

$$\widehat{A}_{22} = A_{22} - A_{21}A_{11}^{-1}A_{12} = A_{22} - L_{21}U_{12}$$

being the Schur complement. We call $p > 0$ the **pivot-block size**, namely the number of row/column pivots performed before each update to the Schur sketching matrix Ψ. With the notation established above, we have $\widehat{A}_{22} = S_{n-m+p+1}$.

Step 1 of Algorithm 3.1 requires that we perform column pivoting on $\widehat{A}_{22}$. To do this work, we need to multiply $\widehat{A}_{22}$ by a random matrix. Instead of generating a new random matrix, we introduce a simple and efficient procedure, **Successive Schur Sketching (SSS)**. We partition Ω and Ψ accordingly as $\Omega = \begin{pmatrix} \Omega_P & \Omega_R \end{pmatrix}$ and

$$\begin{aligned} \Psi = \begin{pmatrix} \Psi_P & \Psi_R \end{pmatrix} &= \begin{pmatrix} \Omega_P & \Omega_R \end{pmatrix} \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \\ &= \begin{pmatrix} \Omega_P A_{11} + \Omega_R A_{21} & \Omega_P A_{12} + \Omega_R A_{22} \end{pmatrix}. \end{aligned}$$

Now we compute the next Schur sketching matrix $\widehat{\Psi}$ for the Schur complement $\widehat{A}_{22}$ as

$$\begin{aligned} \widehat{\Psi} &\stackrel{def}{=} \Omega_R \widehat{A}_{22} = \Omega_R\left(A_{22} - L_{21}U_{12}\right) = \left(\Omega_P A_{12} + \Omega_R A_{22}\right) - \left(\Omega_P A_{12} + \Omega_R L_{21}U_{12}\right) \\ &= \Psi_R - \left(\Omega_P L_{11} + \Omega_R L_{21}\right) U_{12}. \end{aligned} \tag{3.2}$$

Thus, in SSS we use Ω_R, a submatrix of Ω, for the new sample matrix $\widehat{\Psi}$. If done directly, it would take $2r(m-p)^2$ flops to compute $\widehat{\Psi}$ as a matrix-matrix product $\Omega_R\widehat{A}_{22}$. However, since Ψ_R is part

of Ψ and was computed in the previous steps, $\widehat{\Psi}$ can instead be computed via equation (3.2) in about $4rmp$ flops, a very large savings for $p \ll m$. Indeed, this random matrix reuse will prove critically important in reducing the overall cost of computing all sample matrices by Algorithm GERCP. Later on we will further show that Algorithm GERCP will be as reliable as sampling the Schur matrices $\widehat{A}_{22}$ without random matrix reuse.

Given Ψ, the new sample matrix $\widehat{\Psi}$ can be updated in about $4rpm$ flops with the above formula, which is much cheaper than the $O(rn^2)$ flops required for a direct computation. It costs another $2rm$ flops to select a column with sufficiently large column norm. If we perform a column pivot once every block elimination step, the total overhead due to column pivoting includes the computation of the initial sampling matrix, its update at every column pivot, and column selection based on the column norms of the updated sample matrix. These costs add up to

$$4rn^2 + \sum_{k=1}^{n/p} (4rp + 2r)\,(n - (k-1)p) \approx (6r + r/p)\,n^2$$

flops, which is much smaller than the $O(n^3)$ additional flops required by GE2CP. One can preform a floating point error analysis to show that successively updating Ψ at each stage via equation (3.2) as follows.

LEMMA 3.1. *Suppose equation* (3.2) *is continuously used to produce each sketching matrix* Ψ_k. *Then, each sketching matrix in floating point arithmetic is close to the corresponding sketching matrix in real arithmetic as given by*

$$\begin{aligned}
&\left|fl\big(\widehat{\Psi}_{k+1}\big) - \widehat{\Psi}_{k+1}\right| \\
&\quad \le \epsilon_{mach}\,(1+\epsilon_{mach})^k \left(\sum_{i=1}^{k} \left|\Psi_i\,(:,(k+1)p:n)\right| + 5\,|\Omega|\,\left|L_{(k+1)p}\right|\,\left|U_{(k+1)p}(:,(k+1)p:n)\right| \right)
\end{aligned}$$

REMARK 3.2. *Later in the paper, we will provide a probabilistic analysis of this lemma to get rid of the randomness originating from the sampling matrix* Ω.

Proof. A brief floating point analysis proceeds as follows

$$\begin{aligned}
\left|fl\big((\Omega_P L_{11} + \Omega_R L_{21})U_{12}\big) - (\Omega_P L_{11} + \Omega_R L_{21})U_{12}\right| &\le 4\epsilon_{mach}(|\Omega_P|\,|L_{11}| + |\Omega_R|\,|L_{21}|)\,|U_{12}| \\
&= 4\epsilon_{mach} \begin{pmatrix} |\Omega_P| & |\Omega_R| \end{pmatrix} \begin{pmatrix} |L_{11}| \\ |L_{21}| \end{pmatrix} |U_{12}|
\end{aligned}$$

Place this together with equation (3.2) in floating point arithmetic to arrive at

$$\begin{aligned}
&\left|fl\big(\widehat{\Psi}_{k+1}\big) - \widehat{\Psi}_{k+1}\right| \\
&\quad \le \underbrace{\left|fl\big(\Psi_{k,R}\big) - \Psi_{k,R}\right|}_{\text{past update errors}} (1+\epsilon_{mach}) + \epsilon_{mach}\,|\Psi_{k,R}| + 5\epsilon_{mach} \begin{pmatrix} |\Omega_P| & |\Omega_R| \end{pmatrix} \begin{pmatrix} |L_{11}| \\ |L_{21}| \end{pmatrix} |U_{12}|
\end{aligned}$$

Applying simple induction argument to the inequality above gives the desired result. ∎

REMARK 3.3. *This rounding error analysis applies equally well to the* L *and* U *factors computed in finite precision arithmetic. We did not make the distinction to avoid introducing yet more notation.* When U_{11} is well-conditioned, $\widehat{\Psi}$ can be updated by the more efficient formula

$$\begin{aligned}
\widehat{\Psi} &\overset{def}{=} \Omega_R \widehat{A}_{22} = \Omega_R\,(A_{22} - L_{21}U_{12}) = (\Omega_P A_{12} + \Omega_R A_{22}) - (\Omega_P A_{11} + \Omega_R A_{21})\,A_{11}^{-1} A_{12} \\
&= \Psi_R - \left(\Psi_P U_{11}^{-1}\right) U_{12}, \qquad (3.3)
\end{aligned}$$

which costs about $2rpm + 2rm$ flops.

3.2.2. Column pivot quality and column growth factor for GERCP. We present classical Gaussian elimination with randomized complete pivoting (GERCP) below as Algorithm 3.2. Provided that the sampling dimension r is large enough, in this section, we show that, with high probability, each GERCP pivot column has an ℓ_2 length within a constant factor of the largest column, i.e. the GE2CP pivot column. Using this property, we then prove a sub exponential upper bound on the column growth factor for GERCP that holds with probability not less than $1-\delta$ for any user-defined quantity $\delta > 0$. For the sake of simplicity and theoretical guarantees, we will focus on the case with pivot-block size $p = 1$ for the rest of the paper.

ALGORITHM 3.2. **Gaussian Elimination with Randomized Complete Pivoting**

Input: *$n \times n$ matrix A, sampling dimension $r > 0$*
Output: *lower triangular L, upper triangular U, row permutation Π_r, column permutation Π_c.*

---

sample $\Omega(i,j) \sim \mathcal{N}(0,1)$ *for all* $1 \le i \le r$ *and* $1 \le j \le n$
compute $\Psi = \Omega A$
for $k = 1, \cdots, n-1$ **do**

1. **compute** $\alpha = \underset{k \le j \le n}{\arg\max} \|\Psi(:,j)\|_2$.
2. **swap** *columns k and α of A and Ψ.*
3. **compute** $\beta = \underset{k \le i \le n}{\arg\max} |A(i,k)|$.
4. **swap** *rows k and β of A.*
5. **compute** $A(k+1:n,k) = A(k+1:n,k)/A(k,k)$;
6. **compute** $A(k+1:n,k+1:n) = A(k+1:n,k+1:n) - A(k+1:n,k) * A(k,k+1:n)$;
7. **update** $\Psi(:,k:n)$ *with either* (3.2) *or* (3.3) *in real arithmetic*

For ease of notation, we consider a scalar version of GERCP where column pivoting and row pivoting are done one column/one row at a time. The column pivots Π_c are chosen from randomized column pivoting. The row pivots Π_r are also random, but if we fix the column pivots Π_c then the row pivots are deterministic and are uniquely determined by the top-heavy property (2.6). Let $S_k^{\Pi_c} \in \mathbb{R}^{(k:n)\times(k:n)}$ be the Schur complement after the first $k-1$ steps of column and row pivoting and Gaussian elimination, and let $\Omega_k^{\Pi_c} = \Omega(:, \pi_r(k:n)) \in \mathbb{R}^{r\times(n-k+1)}$ be the submatrix of Ω whose columns correspond to the rows of $S_k^{\Pi_c}$. We now define the following events for $1 \le k \le i \le n$:

$$\overline{\mathbf{C}}_{i,k}^{\Pi_c} = \left\{ \left\| \frac{1}{\sqrt{r}} \Omega_k^{\Pi_c} S_k^{\Pi_c}(:,i) \right\|_2 \le \sqrt{1+\epsilon} \left\| S_k^{\Pi_c}(:,i) \right\|_2 \right\} \tag{3.4}$$

$$\underline{\mathbf{C}}_{i,k}^{\Pi_c} = \left\{ \sqrt{1-\epsilon} \left\| S_k^{\Pi_c}(:,i) \right\|_2 \le \left\| \frac{1}{\sqrt{r}} \Omega_k^{\Pi_c} S_k^{\Pi_c}(:,i) \right\|_2 \right\} \tag{3.5}$$

$$\mathbf{C}_{i,k}^{\Pi_c} = \left\{ \sqrt{1-\epsilon} \left\| S_k^{\Pi_c}(:,i) \right\|_2 \le \left\| \frac{1}{\sqrt{r}} \Omega_k^{\Pi_c} S_k^{\Pi_c}(:,i) \right\|_2 \le \sqrt{1+\epsilon} \left\| S_k^{\Pi_c}(:,i) \right\|_2 \right\} \tag{3.6}$$

$$= \overline{\mathbf{C}}_{i,k}^{\Pi_c} \bigcap \underline{\mathbf{C}}_{i,k}^{\Pi_c} \tag{3.7}$$

By Definition 2.9, $\mathbf{C}_{i,k}^{\Pi_c}$ describes the event where the i^{th} column of the k^{th} Schur complement $S_k^{\Pi_c}$ satisfies the $\epsilon-$JL condition under random mapping $\Omega_k^{\Pi_c}$. We also define for each k

$$\mathbf{D}_{k,k}^{\Pi_c} = \left\{ \left\| \Omega_k^{\Pi_c} S_k^{\Pi_c}(:,k) \right\|_2 \ge g \left\| \Omega_k^{\Pi_c} S_k^{\Pi_c} \right\|_{1,2} \right\}.$$

which describes the event where no column is swapped under threshold randomized column pivoting

at step k for some fixed threshold $0 < g \geq 1$. We further define for $k < i$ and $\gamma_k \overset{def}{=} \left\|\Omega_k^{\Pi_c} S_k^{\Pi_c}\right\|_{1,2}$,

$$\mathbf{D}_{i,k}^{\Pi_c} = \left\{ \left\|\Omega_k^{\Pi_c} S_k^{\Pi_c}(:,k)\right\|_2 < g\gamma_k,\ \left\|\Omega_k^{\Pi_c} S_k^{\Pi_c}(:,j)\right\|_2 < \gamma_k,\ k < j \leq i,\ \left\|\Omega_k^{\Pi_c} S_k^{\Pi_c}(:,i)\right\|_2 = \gamma_k. \right\}$$

Thus, $\mathbf{D}_{i,k}^{\Pi_c}$ is the event where columns k and i are swapped. Thus the event

$$\mathcal{D}^{(\Pi_c)} \overset{def}{=} \bigcap_{1 \leq k \leq n} \mathbf{D}_{\alpha_k,k}^{\Pi_c} \tag{3.8a}$$

uniquely defines the column permutation Π_c and by extension, row permutation Π_r, and the event

$$\mathcal{C}^{(\Pi_c)} \overset{def}{=} \bigcap_{1 \leq k \leq n} \left(\left(\bigcap_{\substack{k \leq j \leq n \\ j \neq \alpha_k}} \underline{\mathbf{C}}_{j,k}^{\Pi_c} \right) \bigcap \overline{\mathbf{C}}_{\alpha_k,k}^{\Pi_c} \right) \tag{3.8b}$$

defines the Gaussian elimination process where every column in every Schur complement satisfies the $\epsilon-$JL condition for the column permutation Π_c; and the event

$$\mathcal{C} \overset{def}{=} \bigcup_{\Pi_c} \left(\mathcal{C}^{(\Pi_c)} \bigcap \mathcal{D}\left(\Pi_c\right) \right) \tag{3.8c}$$

defines set of Gaussian elimination processes where every column in each Schur complement produced by the algorithm satisfies the $\epsilon-$JL condition during the factorization process. Note that event $\mathcal{C}$ describes the randomized Gaussian elimination process, whereas event $\mathcal{C}^{(\Pi_c)}$ describes a particular incidence of $\mathcal{C}$ conditional on a particular column permutation Π_c.

Since GERCP performs partial pivoting at every step of elimination, it will successfully compute an LU factorization of any given matrix. Additionally, the events $\mathcal{D}^{(\Pi_c)}$ over the set of permutations are mutually exclusive by definition. In other words,

$$\mathbb{P}\left(\bigcup_{\Pi_c} \mathcal{D}^{(\Pi_c)} \right) = \sum_{\Pi_c} \mathbb{P}\left(\mathcal{D}^{(\Pi_c)} \right) = 1, \tag{3.9}$$

where the set union is over all possible permutations Π_c. Lemma 3.2 below shows that with large probability, in GERCP every column in every Schur complement satisfies the $\epsilon-$JL condition during the elimination process regardless of which permutation Π_c is chosen.

LEMMA 3.2 (Randomized norm preservation). *Given $\epsilon, \delta \in (0,1)$ and $g \in (0,1]$. Let Ψ be defined by equation* (3.1). *Choose $r > \dfrac{4}{\epsilon^2 - \epsilon^3} \ln\left(\dfrac{n(n+1)}{2\delta}\right)$. Then we must have*

$$\|S_k\|_{1,2} \leq \frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\,\|S_k(:,\alpha_k)\|_2 \tag{3.10}$$

for all $1 \leq k \leq n$ with probability no less than $1 - \delta$.

REMARK 3.4. *This lemma tells us that the factor $\sqrt{\frac{1-\epsilon}{1+\epsilon}}$ plays a similar role to the* **column thresholding factor** *g as the two factors appear multiplied together in the same part of the above inequality. Therefore, it is fruitful to view and label the term $\sqrt{\frac{1-\epsilon}{1+\epsilon}}$ from Johnson-Lindenstrauss as the* **artificial column thresholding factor**.

Proof. By equations (3.8), the event

$$\bigcup_{\Pi_c}\left(\mathcal{C}^{(\Pi_c)}\bigcap\mathcal{D}\left(\Pi_c\right)\right)\subseteq\left\{\left\|S_k\right\|_{1,2}\leq\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\left\|S_k(:,\alpha_k)\right\|_2,\ \text{for all } 1\leq k\leq n\right\}$$

defines a subset of the set of outcomes that satisfies our desired result. To show this take any fixed choice of column pivots Π_c. It is trivial to get $\|S_k(:,\alpha_k)\|\leq\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\left\|S_k(:,\alpha_k)\right\|_2$ without any probability. If our outcome is in both $\mathcal{C}\left(\Pi_c\right)$ and $\mathcal{D}\left(\Pi_c\right)$, then for any $1\leq k\leq j\leq n$ and $j\neq\alpha_k$

$$\begin{aligned}
(1-\epsilon)\left\|S_k(:,j)\right\|_2^2&\leq\left\|\frac{1}{\sqrt{r}}\Psi_k(:,j)\right\|_2^2 && \text{(by event } \underline{\mathbf{C}}_{j,k}^{\Pi_c})\\
&\leq\frac{1}{g^2}\left\|\frac{1}{\sqrt{r}}\Psi_k(:,\alpha_k)\right\|_2^2 && \text{(by event } \mathcal{D}(\Pi_c))\\
&\leq\frac{1}{g^2}(1+\epsilon)\left\|S_k(:,\alpha_k)\right\|_2^2 && \text{(by event } \overline{\mathbf{C}}_{\alpha_k,k}^{\Pi_c})
\end{aligned}$$

Choosing $j=\arg\max_{k\leq j\leq n}\left\|S_k(:,j)\right\|_2$ gives us line (3.10) and our desired set containment. Next, we must bound the probability of success from below. It follows from the definition of conditional probability that

$$\begin{aligned}
\mathbb{P}\left(\bigcup_{\Pi_c}\left(\mathcal{C}\left(\Pi_c\right)\cap\mathcal{D}\left(\Pi_c\right)\right)\right)&=\sum_{\Pi_c}\mathbb{P}\left(\mathcal{C}\left(\Pi_c\right)\cap\mathcal{D}\left(\Pi_c\right)\right)\\
&=\sum_{\Pi_c}\mathbb{P}\left(\mathcal{C}\left(\Pi_c\right)|\mathcal{D}\left(\Pi_c\right)\right)\mathbb{P}\left(\mathcal{D}\left(\Pi_c\right)\right) \qquad (3.11)
\end{aligned}$$

Below we derive a lower bound on the right hand side of equation (3.11). Consider any given event $\mathcal{D}\left(\Pi_c\right)$ for which $\mathbb{P}\left(\mathcal{D}\left(\Pi_c\right)\right)>0$. This implies that the permutation Π_c is given. As before, for each $1\leq k\leq n-1$, let $S_k^{\Pi_c}\in\mathbb{R}^{(k:n)\times(k:n)}$ be the Schur complement after the first $k-1$ steps of column and row pivoting and Gaussian elimination, and let $\Omega_k^{\Pi_c}=\Omega(:,\Pi_r(k:n))\in\mathbb{R}^{r\times(n-k+1)}$ be the submatrix of Ω whose columns correspond to the rows of $S_k^{\Pi_c}$. With this notation, we can write

$$\begin{aligned}
\mathbb{P}\left(\mathcal{C}\left(\Pi_c\right)|\mathcal{D}\left(\Pi_c\right)\right)&=\mathbb{P}\left(\bigcap_{1\leq k\leq n}\left(\left(\bigcap_{\substack{k\leq j\leq n\\ j\neq\alpha_k}}\underline{\mathbf{C}}_{j,k}^{\Pi_c}\right)\bigcap\overline{\mathbf{C}}_{\alpha_k,k}^{\Pi_c}\right)\right)\\
&=1-\mathbb{P}\left(\bigcup_{1\leq k\leq n}\left(\left(\bigcup_{\substack{k\leq j\leq n\\ j\neq\alpha_k}}\left(\underline{\mathbf{C}}_{j,k}^{\Pi_c}\right)^c\right)\bigcup\left(\overline{\mathbf{C}}_{\alpha_k,k}^{\Pi_c}\right)^c\right)\right)\\
&\geq 1-\sum_{k=1}^{n}\left(\sum_{\substack{j=k\\ j\neq\alpha_k}}^{n}\mathbb{P}\left(\left(\underline{\mathbf{C}}_{j,k}^{\Pi_c}\right)^c\right)+\mathbb{P}\left(\left(\overline{\mathbf{C}}_{\alpha_k,k}^{\Pi_c}\right)^c\right)\right)\\
&\geq 1-\sum_{k=1}^{n}(n-k+1)\exp\left(-\frac{(\epsilon^2-\epsilon^3)r}{4}\right)\\
&=1-\frac{n(n+1)}{2}\exp\left(-\frac{(\epsilon^2-\epsilon^3)r}{4}\right)
\end{aligned}$$

where the third line comes from Lemma 2.6 and the fourth line is from the application of Lemma 2.8. Plugging this into line (3.11), we have that

$$\mathbb{P}\left(\bigcup_{\Pi_c}\left(\mathcal{C}\left(\Pi_c\right)\cap\mathcal{D}\left(\Pi_c\right)\right)\right) \geq \left(1-\frac{n(n+1)}{2}\exp\left(-\frac{(\epsilon^2-\epsilon^3)r}{4}\right)\right)\sum_{\Pi_c}\mathbb{P}\left(\mathcal{D}\left(\Pi_c\right)\right)$$
$$= 1-\frac{n(n+1)}{2}\exp\left(-\frac{(\epsilon^2-\epsilon^3)r}{4}\right)$$

where the last line is achieved by line (3.9). In order to bound the last line from below by $1-\delta$, we require the $r > \frac{4}{\epsilon^2-\epsilon^3}\ln\left(\frac{n(n+1)}{2\delta}\right)$. □

Below is our main theoretical result on column growth upper bound for GERCP.

THEOREM 3.3 (Column Growth Factor for GERCP). *Choose $\epsilon,\delta\in(0,1)$ and $g\in(0,1]$. If $r > \frac{4}{\epsilon^2-\epsilon^3}\ln\left(\frac{n(n+1)}{2\delta}\right)$, then the pivot growth factor of Algorithm 1.2 satisfies*

$$\rho_{col}^{gercp}(\mathbf{A}) \leq \frac{1}{g^2}\frac{1+\epsilon}{1-\epsilon}\sqrt{e(n+1)}n^{1+\ln\left(g\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right)n^{\frac{1}{2}\ln(n)}}$$

with probability greater than $1-\delta$, otherwise

$$\rho_{col}^{gercp}(\mathbf{A}) \leq \frac{1}{\sqrt{n}}2^{n-1}$$

REMARK 3.5. *What happens if we are extremely unlucky? There is a chance, not exceeding the user-chosen positive probability δ, that at least one column in some Schur complement is not well-preserved. In this case, GERCP could end up picking the wrong column pivot. This, however, is far from a disaster as we still perform the deterministic partial pivoting at every step. Therefore, we can view these randomized column swaps as an "insurance policy" against large element growth because with high probability we will attain the growth bounds in Theorem 3.3, but we will definitely attain GEPP growth bounds.*

Proof. At the k^{th} stage of randomized complete pivoting, our column pivot choice is given by

$$\ell_k = \underset{k\leq j\leq n}{\arg\max}\,\|\Psi_k(:,j)\|_2^2$$
$$\alpha_k = \begin{cases} \ell_k & \text{if } g\,\|\Psi_k(:,\alpha_k)\|_2 \geq \|\Psi_k(:,k)\|_2 \\ k & \text{otherwise}\end{cases}$$

Then the row pivot at the k^{th} stage is given as β_k where

$$\beta_k = \underset{k\leq i\leq n}{\arg\max}\,|S_k(i,\alpha_k)| \tag{3.12}$$

At this point, we proceed in a fashion similar to Wilkinson's element growth proof for complete pivoting [28]. Let $p_k = |S_k(\beta_k,\alpha_k)|$ be the modulus of the pivot element of S_k. Also, let $c_k = \|S_k(:,\alpha_k)\|_2$ be the ℓ_2-norm of the pivot column of S_k. It is important to note that (3.12) along with Lemma 2.1 implies that $c_k \leq \sqrt{n-k+1}p_k$. (i.e. The last line works because $p_k = U(k,k)$ from our a priori pivoting so p_k is an entry of $S_k^{\Pi_c}$.) Then, we have that

$$\left|\det\left(S_k^{\Pi_c}(k:m,k:m)\right)\right| = \prod_{j=k}^{m}p_j \geq \prod_{j=k}^{m}\frac{1}{\sqrt{n-j+1}}c_j \tag{3.13}$$

which holds from the LU decomposition of $S_k^{\Pi_c}$ since the L factor is unit diagonal and the U factor has the p_j's as its diagonal. We can also apply Theorem 2.5 (Hadamard's Inequality) to the determinant of $S_k^{\Pi_c}(k:m, k:m)$ to get

$$\begin{aligned}\left|\det\left(S_k^{\Pi_c}(k:m,k:m)\right)\right| &\leq \prod_{j=k}^{m}\left\|S_k^{\Pi_c}(k:m,j)\right\|_2 \leq \prod_{j=k}^{m}\left\|S_k^{\Pi_c}(:,j)\right\|_2 \\ &\leq \left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right)^{m-k}\|S_k(:,\alpha_k)\|_2^{m-k+1} = \left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right)^{m-k} c_k^{m-k+1} \quad (3.14)\end{aligned}$$

where the last expression of the first line is achieved by taking the ℓ_2 norm of the entire j^{th} column instead of the first few entries of the j^{th} column. Also, the second line of the above is achieved from Theorem 3.2 by applying (3.10) for each $j \neq \alpha_k$. Combining our inequalities (3.13) and (3.14) for $|\det(S_k(k:m,k:m))|$, we get

$$\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right)^{m-k}\sqrt{\frac{(n-k+1)!}{(n-m)!}}c_k^{m-k+1} \geq \prod_{j=k}^{m} c_j$$

for all $1 \leq k \leq m \leq n$. Define $q_k = \ln(c_k)$. Canceling one c_k on both sides and taking logarithms on both sides, we get

$$(m-k)\ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right) + \sum_{j=k}^{m}\ln\sqrt{n-j+1} + (m-k)\,q_k \geq \sum_{j=k+1}^{m} q_j$$

Dividing by $m-k$, moving each term with any q_j for $k \leq j < m$ to one side and everything else to the other,

$$q_k - \frac{1}{m-k}\sum_{j=k+1}^{m-1} q_j \geq \frac{1}{m-k}q_m - \frac{1}{m-k}\sum_{j=k}^{m}\ln\sqrt{n-j+1} - \ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right)$$

Next, we combine all the inequalities for each k between $1 \leq k < m$ to get

$$\begin{pmatrix} 1 & -\frac{1}{m-1} & -\frac{1}{m-1} & \cdots & -\frac{1}{m-1} & -\frac{1}{m-1} \\ 0 & 1 & -\frac{1}{m-2} & \cdots & -\frac{1}{m-2} & -\frac{1}{m-2} \\ 0 & 0 & 1 & \cdots & -\frac{1}{m-3} & -\frac{1}{m-3} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & -\frac{1}{2} \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix}\begin{pmatrix} q_1 \\ q_2 \\ q_3 \\ \vdots \\ q_{m-2} \\ q_{m-1} \end{pmatrix} \geq \begin{pmatrix} \frac{1}{m-1}q_m \\ \frac{1}{m-2}q_m \\ \frac{1}{m-3}q_m \\ \vdots \\ \frac{1}{2}q_m \\ q_m \end{pmatrix} - \begin{pmatrix} \frac{1}{m-1}\ln\sqrt{\frac{n!}{(n-m)!}} \\ \frac{1}{m-2}\ln\sqrt{\frac{(n-1)!}{(n-m)!}} \\ \frac{1}{m-3}\ln\sqrt{\frac{(n-2)!}{(n-m)!}} \\ \vdots \\ \frac{1}{2}\ln\sqrt{\frac{(n-m+3)!}{(n-m)!}} \\ \ln\sqrt{\frac{(n-m+2)!}{(n-m)!}} \end{pmatrix} - \begin{pmatrix} \ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right) \\ \ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right) \\ \ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right) \\ \vdots \\ \ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right) \\ \ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right) \end{pmatrix}$$

and we express the above matrix inequality as

$$B \qquad \mathbf{q} \quad \overset{def}{\geq} \quad \mathbf{v}_1 \quad - \quad \mathbf{v}_2 \quad - \quad \mathbf{v}_3$$

Lemma 5.2 gives us a closed form expression for B^{-1}, which happens to be a non-negative matrix. Since B^{-1} only has non-negative entries, the vector inequality above is preserved after multiplying

B^{-1} on both sides. To complete this proof, we only need the top row of this vector inequality. Since q_1 is the first entry of $\mathbf{q}$, we also apply $\mathbf{e}_1^T$ to both sides of the inequality to reduce it to a scalar inequality

$$q_1 = \mathbf{e}_1^T \mathbf{q} \geq \mathbf{e}_1^T B^{-1} \mathbf{v}_1 - \mathbf{e}_1^T B^{-1} \mathbf{v}_2 - \mathbf{e}_1^T B^{-1} \mathbf{v}_3 \tag{3.15}$$

Lemma 5.2 also tells us that the first row of B^{-1} is

$$\mathbf{e_1}^T B^{-1} = \begin{bmatrix} 1 & \frac{1}{m-1} & \frac{1}{m-2} & \cdots & \frac{1}{3} & \frac{1}{2} \end{bmatrix}$$

which we now use to compute/bound $\mathbf{e}_1^T B^{-1} \mathbf{v}_1$, $\mathbf{e}_1^T B^{-1} \mathbf{v}_2$ and $\mathbf{e}_1^T B^{-1} \mathbf{v}_3$:

$$\mathbf{e}_1^T B^{-1} \mathbf{v}_1 = q_m \left(\frac{1}{m-1} + \sum_{j=1}^{m-2} \frac{1}{(j+1)j} \right) = q_m$$

where the last equality was achieved by Lemma 5.1. Next, we compute $\mathbf{e}_1^T B^{-1} \mathbf{v}_2$ which we call the Wilkinson term

$$\begin{aligned}
\mathbf{e}_1^T B^{-1} \mathbf{v}_2 &= \frac{1}{m-1} \sum_{j=1}^{m} \ln \sqrt{n-m+j} + \sum_{k=1}^{m-2} \frac{1}{(k+1)k} \sum_{j=1}^{k+1} \ln \sqrt{n-m+j} \\
&= \frac{1}{m-1} \sum_{j=1}^{m} \ln \sqrt{n-m+j} + \sum_{j=1}^{m-1} \sum_{k=\max\{j-1,1\}}^{m-2} \frac{1}{(k+1)k} \ln \sqrt{n-m+j} \\
&= \frac{1}{m-1} \ln \sqrt{n} + \sum_{j=1}^{m-1} \left(\frac{1}{m-1} + \sum_{k=\max\{j-1,1\}}^{m-2} \frac{1}{(k+1)k} \right) \ln \sqrt{n-m+j} \\
&= \frac{1}{m-1} \ln \sqrt{n} + \sum_{j=1}^{m-1} \frac{1}{\max\{j-1,1\}} \ln \sqrt{n-m+j} \qquad \text{(Lemma 5.1)} \\
&= \ln \sqrt{n-m+1} + \sum_{j=2}^{m} \frac{1}{j-1} \ln \sqrt{n-m+j} \\
&\stackrel{def}{=} \ln \sqrt{n-m+1} + \ln f(m, n-m)
\end{aligned}$$

where we define the *generalized Wilkinson function* to be

$$f(m,t) \stackrel{def}{=} \sqrt{(2+t)^1 (3+t)^{\frac{1}{2}} (4+t)^{\frac{1}{3}} \cdots (m+t)^{\frac{1}{m-1}}}$$

Next comes $\mathbf{e}_1^T B^{-1} \mathbf{v}_3$ or the thresholding term

$$\mathbf{e}_1^T B^{-1} \mathbf{v}_3 = \ln \left(\frac{1}{g} \sqrt{\frac{1+\epsilon}{1-\epsilon}} \right) \sum_{j=1}^{m-1} \frac{1}{j} \leq (1 + \ln(m-1)) \ln \left(\frac{1}{g} \sqrt{\frac{1+\epsilon}{1-\epsilon}} \right)$$

Plugging all of this into (3.15), we have

$$q_1 \geq q_m - \ln \sqrt{n-m+1} - \ln f(m, n-m) - (1 + \ln(m-1)) \ln \left(\frac{1}{g} \sqrt{\frac{1+\epsilon}{1-\epsilon}} \right)$$

Taking the exponential of both sides and rearranging terms, we have that for all $1 \leq m \leq n$

$$\frac{c_m}{c_1} \leq \sqrt{n-m+1} f(m, n-m) \frac{1}{g} \sqrt{\frac{1+\epsilon}{1-\epsilon}} (m-1)^{\ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right)} \tag{3.16}$$

Next, we need to relate the ratio $\frac{c_k}{c_1}$ to ρ_{col}^{gercp}

$$\begin{aligned}
\rho_{col}^{gercp} &= \frac{\max_k \|S_k\|_{1,2}}{\|S_1\|_{1,2}} \\
&\leq \frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}} \max_k \frac{\|S_k(:,\alpha_k)\|_2}{\|S_1(:,\alpha_1)\|_2} && \text{(Lemma 3.2)} \\
&= \frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}} \max_k \frac{c_k}{c_1} \\
&\leq \frac{1}{g^2}\frac{1+\epsilon}{1-\epsilon} \max_m \sqrt{n-m+1} f(m, n-m)(m-1)^{\ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right)} && \text{(Eqn (3.16))} \\
&\leq \frac{1}{g^2}\frac{1+\epsilon}{1-\epsilon} \max_m \sqrt{e(n-m+2)}(n-m+1)(m-1)^{\ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right)} \\
&\qquad m^{\frac{1}{4}\ln(n)} m^{\frac{1}{4}\ln\left(\frac{n}{m}\right)} (n-m+1)^{\frac{1}{4}\ln(n-m+1)} \\
&\leq \frac{1}{g^2}\frac{1+\epsilon}{1-\epsilon} \max_m \left(\sqrt{e(n-m+2)}(n-m+1)(m-1)^{\ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right)}\right) \\
&\qquad \max_m \left(m^{\frac{1}{4}\ln(n)} m^{\frac{1}{4}\ln\left(\frac{n}{m}\right)}\right) \max_m (n-m+1)^{\frac{1}{4}\ln(n-m+1)} && \text{(Lemma 5.3)} \\
&\leq \frac{1}{g^2}\frac{1+\epsilon}{1-\epsilon} \sqrt{e(n+1)} n^{1+\ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right)} \max_m \left(m^{\frac{1}{4}\ln(n)} m^{\frac{1}{4}\ln\left(\frac{n}{m}\right)}\right) n^{\frac{1}{4}\ln(n)}
\end{aligned}$$

where the third to last line is from Lemma 5.3. Examining the following derivative

$$\frac{d}{dm} \ln\left(m^{\frac{1}{4}\ln(n)} m^{\frac{1}{4}\ln\left(\frac{n}{m}\right)}\right) = \frac{\ln(n)}{m} - \frac{\ln(m)}{m}$$

which equals zero only when $m = n$, and given the concavity, this point attains the max. Thus, we plug this point in to get

$$\rho_{col}^{gercp} \leq \frac{1}{g^2}\frac{1+\epsilon}{1-\epsilon} \sqrt{e(n+1)} n^{1+\ln\left(\frac{1}{g}\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right)} n^{\frac{1}{2}\ln(n)}$$

to get our desired result. □

Next, we slightly add to the statement of the last theorem to add a probabilistic guarantee on the floating point error of the sampling matrix Ψ under the sampling update formula (3.2). The floating error analysis of the sampling matrix must be done under the objective that rounding errors cannot corrupt the column lengths by too much, which is suggested by the following result under small probability of failure. This shows that the floating point error cannot grow quickly or exponentially under the given update scheme.

THEOREM 3.4 (Stability of unconditionally stable sampling update formula). *For Gaussian Elimination with Randomized Complete Pivoting, we have the two guarantees*

$$\rho_{col}^{gercp}(\mathbf{A}) \leq \frac{1}{g^2}\frac{1+\epsilon_{JL}}{1-\epsilon_{JL}} \sqrt{e(n+1)} n^{1+\ln\left(g\sqrt{\frac{1+\epsilon_{JL}}{1-\epsilon_{JL}}}\right)} n^{\frac{1}{2}\ln(n)}$$

and

$$\left\| fl\big(\widehat{\Psi}_{k+1}\big) - \widehat{\Psi}_{k+1} \right\|_{1,2} \leq \epsilon_{mach} \left(1+\epsilon_{mach}\right)^k \sqrt{kr(1+\epsilon_{JL})} \left(\sqrt{k} + 5n\sqrt{n(1+\epsilon_{JL})} \right) \rho_{col}^{gercp}(A) \left\| A \right\|_{1,2}$$

with probability at least $1 - \frac{n(n+1)}{2} \exp\left(-\frac{(\epsilon_{JL}^2 - \epsilon_{JL}^3)r}{4}\right) - \exp\left(-\frac{\epsilon_{JL}^2 rn}{2}\right)$

Proof. First, we apply the relevant norm to both sides to get

$$\begin{aligned} &\left\| fl\big(\widehat{\Psi}_{k+1}\big) - \widehat{\Psi}_{k+1} \right\|_{1,2} \\ &\leq \epsilon_{mach} \left(1+\epsilon_{mach}\right)^k \left(\sum_{i=1}^{k} \left\| \Psi_i\left(:,(k+1):n\right) \right\|_{1,2} + 5 \left\| |\Omega| \, |L_{k+1}| \right\|_2 \left\| |U_{k+1}(:,(k+1):n)| \right\|_{1,2} \right) \\ &\leq \epsilon_{mach} \left(1+\epsilon_{mach}\right)^k \left(\sum_{i=1}^{k} \left\| \Psi_i\left(:,(k+1):n\right) \right\|_{1,2} + 5 \left\| \Omega \right\|_F \left\| L_{k+1} \right\|_F \left\| U_{k+1}(:,(k+1):n) \right\|_{1,2} \right) \end{aligned} \tag{3.17}$$

where we bound the 2-norm of the entrywise absolute value by the frobenius norm to get the last line. Next, Theorem 3.3 and the definition of the event $\overline{\mathbf{C}}_{\alpha_k,k}^{\Pi_c}$ from (3.4) give us both our desired bound on $\rho_{col}^{gercp}(\mathbf{A})$ and

$$\left\| \Psi_k \right\|_{1,2} = \left\| \Psi_k(:,\alpha_k) \right\|_2 \leq \sqrt{r(1+\epsilon_{JL})} \left\| S_k \right\|_{1,2} \leq \sqrt{r(1+\epsilon_{JL})} \rho_{col}^{gercp}(\mathbf{A}) \left\| \mathbf{A} \right\|_{1,2} \tag{3.18}$$

with probability of failure bounded above by $\frac{n(n+1)}{2} \exp\left(-\frac{(\epsilon_{JL}^2 - \epsilon_{JL}^3)r}{4}\right)$. Since GERCP is a top-heavy method from Definition 2.10, we have $\left\| L_{k+1} \right\|_F \leq \sqrt{kn}$ because L_{k+1} is a $n \times k$ matrix with each entry bounded above by 1 in absolute value. Also, Theorem 2.16 gives us that $\left\| U_{k+1} \right\|_{1,2} \leq n\rho_{col}^{gercp}(\mathbf{A}) \left\| \mathbf{A} \right\|_{1,2}$. Finally, we observe that the map $f(\Omega) = \left\| \Omega \right\|_F$ is a lipschitz map with Lipschitz constant $L = 1$ and apply Theorem 5.5 to get

$$\mathbb{P}\left\{ \left\| \Omega \right\|_F \geq \sqrt{rn} + t \right\} \leq e^{-\frac{t^2}{2}}$$

where we applied Jensen's inequality to get $\mathbb{E} \left\| \Omega \right\|_F \leq \sqrt{\mathbb{E} \left\| \Omega \right\|_F^2} = \sqrt{rn}$ from Proposition 10.1 from [13]. Set $t = \epsilon_{JL}\sqrt{rn}$. Then plug this and the inequality (3.18) into the first inequality (3.17) to arrive at our desired conclusion with a union bound. ∎

4. Numerical Experiments.

4.1. Block GERCP. Algorithm 3.2 was presented in a form for ease of presentation of Theorem 3.3. For efficient numerical implementation, we need to develop a block version of GERCP in Algorithm 3.2 to increase locality and memory/cache re-use with more BLAS-3 calls. Algorithm 4.1 below is styled after dgetf2.f and dgetrf.f for block GEPP in LAPACK with double precision floating point numbers. Instead of increasing the pivot-block size from $p = 1$, we introduce a loop-blocksize parameter $b \geq 1$. It repeatedly performs b steps of randomized column pivoting and partial pivoting followed by a blocksize b Schur complement update.

ALGORITHM 4.1. **Block GERCP**

Input: *$n \times n$ matrix A, sampling dimension $r > 0$, block size b*
Output: *lower triangular L, upper triangular U, row permutation Π_r, column permutation Π_c.*

---

sample $\Omega(i,j) \sim \mathcal{N}(0,1)$ *for all* $1 \leq i \leq r$ *and* $1 \leq j \leq n$
compute $\Psi = \Omega A$
for $\underline{k} = 1 : b : n-1$ **do**

set $\overline{k} = \underline{k} + \min\{b, n - \underline{k} + 1\} - 1$
for $k = \underline{k} : \overline{k}$ **do**

1. **compute** $\alpha = \begin{cases} \arg\max_{k \le j \le n} \|\Psi(:, j)\|_2 & \text{if } n - k \ge r \\ \arg\max_{k \le j \le n} \|A(k:n, j)\|_2 & \text{otherwise} \end{cases}$.
2. **swap** *columns* k *and* α *of* A, Ψ *and* Ω *(*).*
3. **compute** $\beta = \arg\max_{k \le i \le n} |A(i, k)|$.
4. **swap** *rows* k *and* β *of* A *(*).*
5. **compute** $A(k{+}1{:}n, k) = A(k{+}1{:}n, k)/A(k, k)$;
6. **compute** $A(k{+}1{:}n, k{+}1{:}\overline{k}) = A(k{+}1{:}n, k{+}1{:}\overline{k}) - A(k{+}1{:}n, k) * A(k, k{+}1{:}\overline{k})$;
7. **compute** $A(k, \overline{k}{+}1{:}n) = A(k, \overline{k}{+}1{:}n) - A(k, \underline{k}{:}k{-}1) * A(\underline{k}{:}k{-}1, \overline{k}{+}1{:}n)$;
8. **update** $\Psi(:, k{:}n)$ *with Algorithm 4.2*

end for
compute $A(\overline{k}{+}1{:}n, \overline{k}{+}1{:}n) = A(\overline{k}{+}1{:}n, \overline{k}{+}1{:}n) - A(\overline{k}{+}1{:}n, \underline{k}{:}\overline{k}) * A(\underline{k}{:}\overline{k}, \overline{k}{+}1{:}n)$;
end for

The main work of Algorithm 4.1 is in the last step, the repeated computation of the matrix $A(\overline{k}{+}1{:}n, \overline{k}{+}1{:}n)$. The outer loop for $\underline{k}$ is similar to the main loop in dgetrf.f, while the inner loop for k is similar to the main loop in dgetf2.f. The main modifications occur on lines 7, 8 and the BLAS-3 Schur complement update after the end of the inner loop. Line 7 updates the U factor so that we can use it to update the sketching matrix on line 8. As in dgetf2.f, the inner loop for k is designed to work on a tall-skinny matrix, whereas the outer loop for $\underline{k}$ performs fast BLAS-3 updates on the rest of the matrix. For practical reasons, we stop using the sampling matrix once the dimensions of the Schur complement become less than or equal to the sampling dimension $r > 0$. It is important to note that the proof of Theorems 3.3 and 3.4 is easily modified to apply to this versions of the algorithm with the exact same guarantees on element growth. We present the procedure for updating the Schur sampling matrix Ψ.

ALGORITHM 4.2. **Update procedure for Schur sampling matrix Ψ**

Input: $r \times n$ *matrix* Ψ, $n \times n$ *working matrix* A, $r \times n$ *random matrix* Ω
Output: $r \times m$ *matrix* Ψ

---

if *pivot* $|A(k, k)| \ge \sqrt{\epsilon_{mach}} \|\Psi_1\|_{1,2}$ **then** *apply Eqn* (3.3) *with*

$$\Psi(:, (k+1):n) \longleftarrow \Psi(:, (k+1):n) - \frac{\Psi(:, k) A(k, (k+1):n)}{A(k, k)}$$

else *apply Eqn* (3.2) *with*

$$\Psi(:, (k+1):n) \longleftarrow \Psi(:, (k+1):n) - [\Omega(:, k) + \Omega(:, (k+1):n) A((k+1):n, k)]\, A(k, (k+1):n)$$

REMARK 4.1. *The first updating formula in Algorithm 4.2 is slightly more efficient than the second one. While we have not observed it in our numerical computations, it potentially could lead to inaccurate column selections for some highly ill-conditioned matrices in pathological cases. Our numerical experiments also suggest that the execution time of Algorithm 4.2 is typically a relatively small fraction of the total execution time of of GERCP. Thus, one might want to use the second updating formula in Algorithm 4.2 for a more robust numerical implementation.*

4.2. Numerical Results. We ran our experiments on two different machines. The runtime results were preformed on a single node of NERSC's Carver machine with two quad-core Intel Xeon

N	3000	5000	7000	9000	11000
t_{rcp} (secs)	2.250	9.623	25.506	53.607	96.662
t_{pp} (secs)	2.006	8.902	24.066	50.934	91.967
$\frac{t_{rcp}-t_{pp}}{t_{pp}}$	12.20%	8.10%	6.00%	5.20%	5.10%

TABLE 4.1
Average run times of GERCP and GEPP over 10 seperate trials.

X5550 2.67 GHz processors and 24 GB of RAM. The rest of the numbers are generated on a laptop with an Intel i7-3632QM CPU and 8GB of RAM. All of the code here was run using the MKL BLAS. We used the open source Netlib version of GEPP. Our version of GERCP was obtained by modifying the Netlib GEPP Fortran code. This allows for an easy and fair comparison between GEPP and GERCP by insuring that the version of GEPP used to compare against has similar cache optimizations. It is worth noting that the Intel MKL version of GEPP is much faster than both Netlib GEPP and GERCP because of superior cache optimizations. Figure 4.1 from our runtime experiments shows that as the matrix size increases, the percent difference in runtime decreases to a negligible amount. This agrees with our theory, which tells us the $O(n^2)$ operations required to maintain the sampling matrix and pivot columns does not grow as quickly as the $O(n^3)$ operations required to actually factor the matrix A. Even when the relative time difference is high, the absolute time difference is unnoticeable for a single factorization as shown in by the run times for $N = 3000$ in Table 4.2 below.

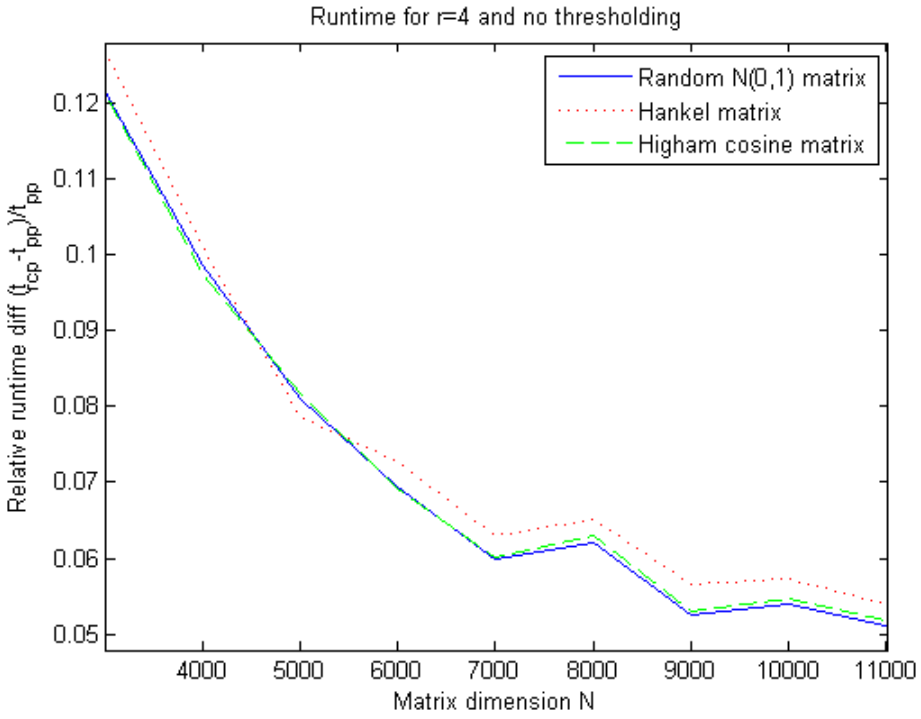

FIG. 4.1. *Comparing the run times of GERCP and GEPP Fortran code each averaged over* 10 *different trials*

In section 2.4.2, we reviewed stability issues associated with to most commonly used Gaussian elimination pivoting strategy, partial pivoting. Now, we produce numerical experiments showing the

improved stability properties of GERCP. The Wilkinson, Generalized Wilkinson and Volterra matrices that we use in Figures 4.2 and 4.3 in these plots are as described in Section 2.4.2. The Wilkinson-type matrices serve as the worst case matrix for GEPP instability with entries that grow exponentially, where the standard Wilkinson matrix has the quickest exponential growth with a base of 2. This is exemplified by the dashed blue and red lines in the log-log plots of element growth and backwards error of GEPP in Figures 4.2 and 4.3. However, GERCP fixes this by impeding element growth with column pivots to leave the backwards error near machine precision. As far as Gaussian elimination goes, the most pathological examples, that we characterize as *passive aggressive* element growth, include the Volterra matrix. These matrices exhibit just enough element growth, but no more, to cause an unacceptable level of backwards error. In contrast, the element/column growth for the Wilkinson-type matrices is so massive that the problem is trivial to detect and fix for any GE algorithm with both row and column pivots to correct. This case too is effortlessly corrected by GERCP as shown by our experiments.

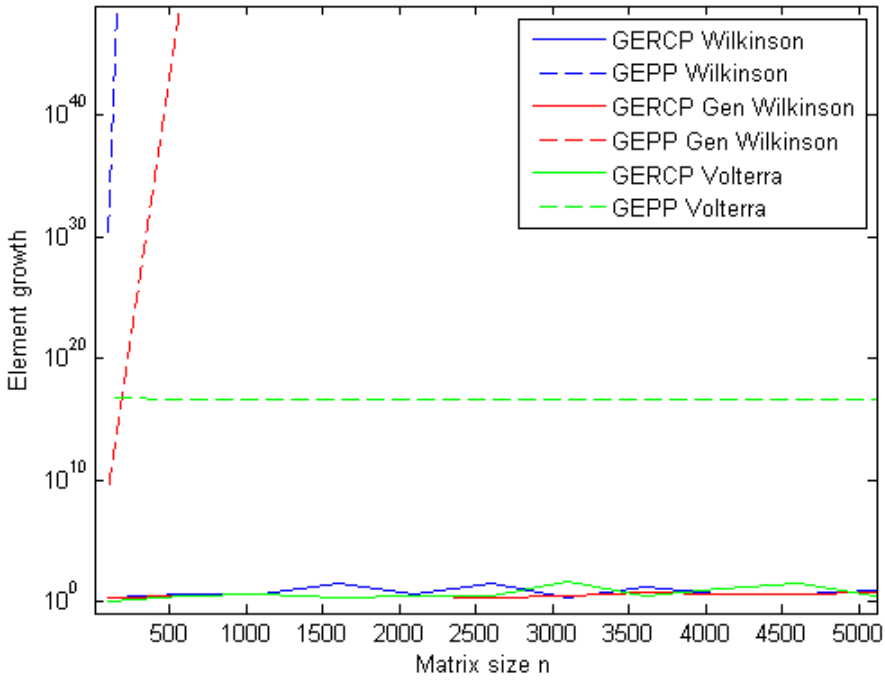

FIG. 4.2. *Element growth for diabolical matrices with GERCP and GEPP Fortran code*

4.3. Backward error for random linear systems. Suppose we wish to solve the linear system $\mathbf{Ax} = \mathbf{b}$ where $\mathbf{A} \in \mathbb{R}^{n\times n}$ iid $\mathcal{N}(0,1)$ standard normal and $\mathbf{b} \in \mathbb{R}^n$ iid $\mathcal{N}(0,1)$ standard normal. This system is known to be well conditioned [3] and the LU factorization is stable under partial pivoting and complete pivoting [25]. We measure the accuracy of a linear solve with the *relative residual*

$$\frac{\|\mathbf{A}\widehat{\mathbf{x}} - \mathbf{b}\|_\infty}{\|\mathbf{A}\|_\infty \|\widehat{\mathbf{x}}\|_\infty}$$

In Figure 4.3, we plot the relative residual error for GERCP with different sampling parameters $r > 0$, along with competing methods like GEPP, GECP, GERP (rook) and GE2CP. While complete

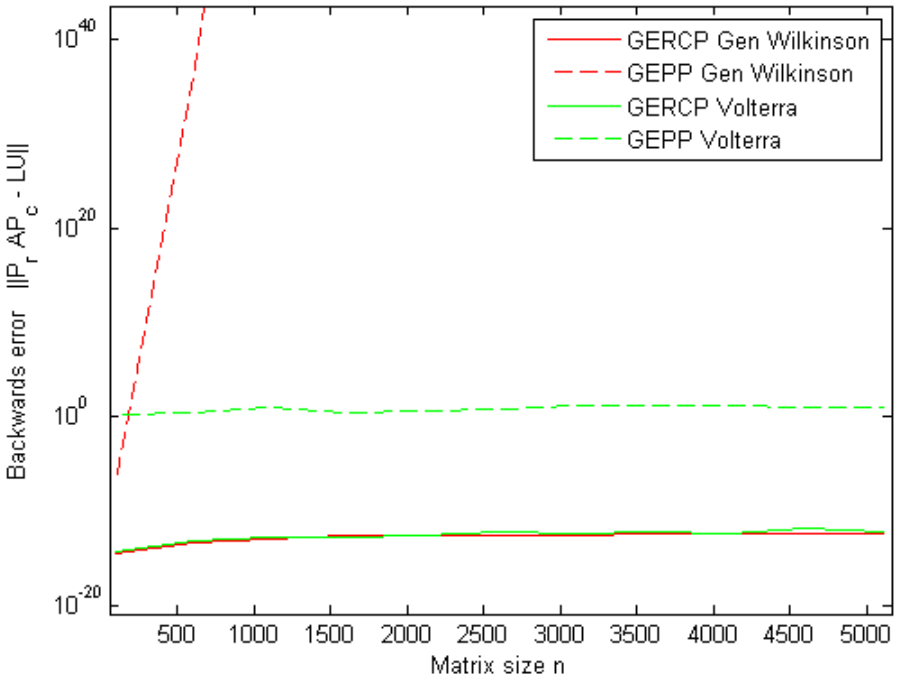

FIG. 4.3. *Backwards error for diabolical matrices with GERCP and GEPP Fortran code*

pivoting consistently obtained the smallest residual, GERCP, GERP and GE2CP all produced similar relative residuals which were clearly better than GEPP and not much worse than that of GECP. This shows that GERCP produces a better quality solution than GEPP even when GEPP is given a well-conditioned system. The Figure 4.3 suggests that the relative residual for the random normal linear system is improved by almost a factor of 2. This savings becomes more important when you work with smaller precession floating point numbers like single or half precession floating point numbers. These smaller precession floating point numbers are becoming commonly used on co-processing platforms like GPGPUs as result in dramatic run time improvements. Also, for different linear systems this improvement in the relative residual can be much higher as in the case of the Wilkinson-type and Volterra matrices.

We also look at the element growth within the LU factors for different pivoting strategies. In [25], Trefethen and Schreiber study the element growth factors for GEPP and GECP on standard normal matrices. They conjecture that $\mathbb{E}\left(\rho_n^{gepp}\right) \approx n^{2/3}$ and $\mathbb{E}\left(\rho_n^{gecp}\right) \approx n^{1/2}$. We plot the element growth for GEPP and GECP along with the element growth for GERP (rook), GE2CP and GERCP with different values of the sampling parameter $r > 0$ in Figure 4.3.

Acknowledgements and Future Work. The authors would like to thank James Demmel and Laura Grigori for many helpful and fruitful discussions. There are many directions for future research into randomized complete pivoting. One direction is to study and provide guarantees for variants of GERCP that allow for a large pivot-block size $p > 1$. The authors believe that such a method could be much faster due to an increase in BLAS-3 operations, and could also provide rank-revealing style guarantees for low-rank approximations generated by incomplete LU factorizations [20]. It will also be important to develop a cache optimized version if this code to be competitive to Intel MKL LAPACK GEPP. Another important avenue for future research is to use these techniques to make

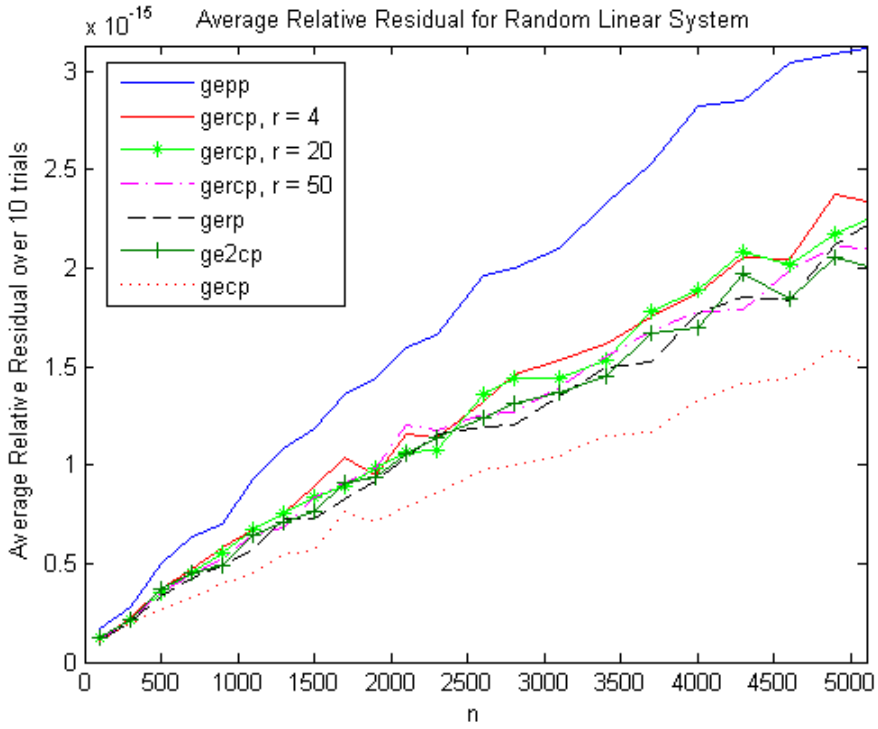

FIG. 4.4. *Average Relative Residual over* 10 *trials. This suggests that GERCP should improve the relative residual of a linear solve over GEPP by at least half the improvement that GECP would provide.*

communication avoiding tournament pivoting LU (CALU) more stable by adding randomized column pivots [5, 12].

REFERENCES

[1] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen. *LAPACK Users' Guide*. SIAM, Philadelphia, PA, second edition, 1994.

[2] V. Bogdanov. *Gaussian Measures*. American Mathematical Society, Providence, RI, 1998.

[3] Z. Chen and J. Dongarra. Condition numbers of Gaussian random matrices. *SIAM J. Matrix Anal. Appl.*, 27:603–620, 2005.

[4] J. Demmel. *Applied Numerical Linear Algebra*. SIAM, Philadelphia, PA, 1997.

[5] Simplice Donfack, Stanimire Tomov, and Jack Dongarra. Dynamically balanced synchronization-avoiding lu factorization with multicore and gpus. In *Parallel & Distributed Processing Symposium Workshops (IPDPSW), 2014 IEEE International*, pages 958–965. IEEE, 2014.

[6] J. Dongarra, J. Bunch, C. Moler, and G. W. Stewart. *Linpack User's Guide*. SIAM, Philadelphia, PA, 1979.

[7] Alan Edelman. The complete pivoting conjecture for gaussian elimination is false. *Mathematica Journal*, 2(2):58–61, 1992.

[8] Leslie V. Foster. Gaussian elimination with partial pivoting can fail in practice. *SIAM Journal on Matrix Analysis and Applications*, 15(4):1354–1362, 1994.

[9] Leslie V. Foster. The growth factor and efficiency of Gaussian elimination with rook pivoting. *Journal of Computational and Applied Mathematics*, 86:177–194, 1997. Corrigendum in JCAM, 98:177, 1998.

[10] Philip E Gill, Walter Murray, and Michael A Saunders. Snopt: An sqp algorithm for large-scale constrained optimization. *SIAM journal on optimization*, 12(4):979–1006, 2002.

[11] Nick Gould. On growth in gaussian elimination with complete pivoting. *SIAM Journal on Matrix Analysis and Applications*, 12(2):354–361, 1991.

[12] Laura Grigori, James W Demmel, and Hua Xiang. Calu: a communication optimal lu factorization algorithm.

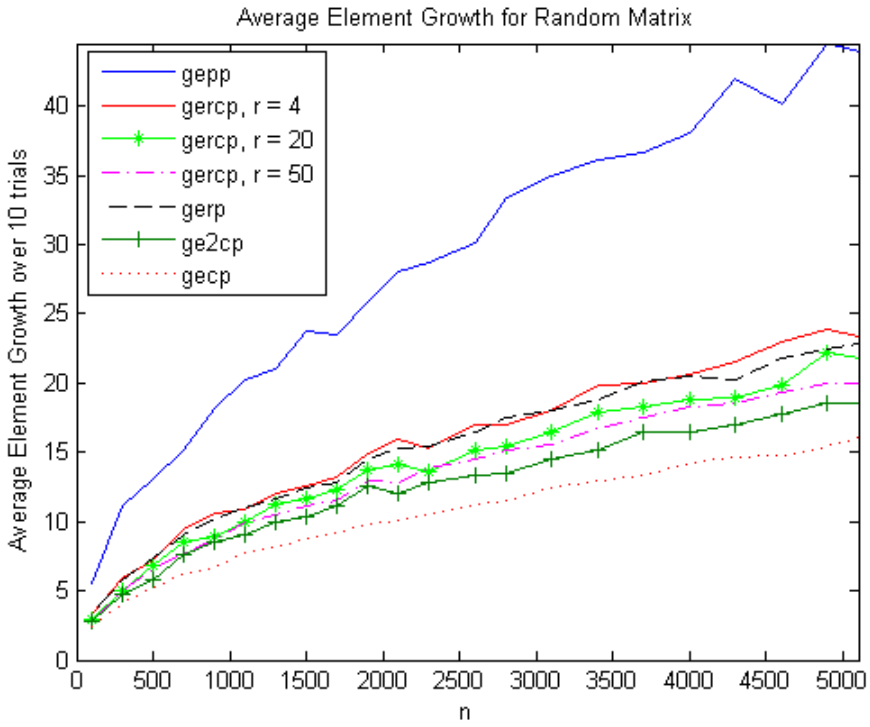

FIG. 4.5. *Average element growth of iid standard Normal random matrix over* 10 *trials.*

SIAM Journal on Matrix Analysis and Applications, 32(4):1317–1350, 2011.

[13] N. Halko, P.-G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53:217–288, 2011.

[14] N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, 1996.

[15] R. A. HORN and C. R. JOHNSON. *Matrix Analysis*. CAMBRIDGE UNIVERSITY PRESS, 1985.

[16] J. M. Hyman and B. Nicolaenko. The Kuramoto-Sivashinsky equation: a bridge between PDE's and dynamical systems. *Physica D*, 18:113–126, 1986.

[17] Fritz John. Extremum problems with inequalities as subsidiary conditions. In *Traces and Emergence of Nonlinear Programming*, pages 197–215. Springer, 2014.

[18] Robert Eugene LaQuey, SM Mahajan, PH Rutherford, and WM Tang. Nonlinear saturation of the trapped-ion mode. *Physical Review Letters*, 34(7):391, 1975.

[19] Lixin Liu and Robert D Russell. Linear system solvers for boundary value odes. *Journal of Computational and Applied Mathematics*, 45(1):103–117, 1993.

[20] L Miranian and Ming Gu. Strong rank revealing lu factorizations. *Linear algebra and its applications*, 367:1–16, 2003.

[21] Larry Neal and George Poole. A geometric analysis of Gaussian elimination. II. *Linear Algebra Appl.*, 173:239–264, 1992.

[22] Victor Y Pan, Guoliang Qian, and Ai-Long Zheng. Randomized preprocessing versus pivoting. *Linear Algebra and Its Applications*, 438(4):1883–1899, 2013.

[23] D Stott Parker. Random butterfly transformations with applications in computational linear algebra. 1995.

[24] George Poole and Larry Neal. The rook's pivoting strategy. *J. Comput. Appl. Math.*, 123(1-2):353–369, 2000. Numerical analysis 2000, Vol. III. Linear algebra.

[25] L. Trefethen and R. Schreiber. Average case analysis of Gaussian elimination. *SIAM J. Mat. Anal. Appl.*, 11(3):335–360, 1990.

[26] Santosh Srinivas Vempala. *The random projection method*, volume 65 of *DIMACS series in discrete mathematics and theoretical computer science*. Providence, R.I. American Mathematical Society, 2004. Appendice p.101-105.

[27] John Von Neumann and Herman H Goldstine. Numerical inverting of matrices of high order. *Bulletin of the*

American Mathematical Society, 53(11):1021–1099, 1947.
[28] J. H. Wilkinson. Error analysis of direct methods of matrix inversion. *J. ACM*, 10:281–330, 1961.
[29] David P Woodruff. Sketching as a tool for numerical linear algebra. *arXiv preprint arXiv:1411.4357*, 2014.
[30] Stephen J. Wright. A collection of problems for which gaussian elimination with partial pivoting is unstable. *SIAM J. Sci. Comput.*, 14(1):231–238, January 1993.

5. Appendix. LEMMA 5.1 (Partial fraction telescoping sum). *Let $q > r > 0$ be positive integers. Then we have*

$$\frac{1}{r} = \frac{1}{q} + \sum_{j=r}^{q-1} \frac{1}{(j+1)j}$$

Proof. Let $s \in \mathbb{N}$ such that $r \leq s < q$. We use common denominators to subtract the two following fractions

$$\frac{1}{s} - \frac{1}{s+1} = \frac{1}{s(s+1)}$$

Then we take the telescoping sum to arrive at our result

$$\begin{aligned} \frac{1}{r} - \frac{1}{q} &= \sum_{s=r}^{q-1} \frac{1}{s} - \frac{1}{s+1} \\ &= \sum_{s=r}^{q-1} \frac{1}{s(s+1)} \end{aligned}$$

□

LEMMA 5.2 (Special Matrix Inverse). *Let $B = (b_{ij})_{1 \leq i,j \leq n} \in \mathbb{R}^{n \times n}$ be an upper triangular matrix given by*

$$b_{ij} = \begin{cases} 0 & \text{if } i > j \\ 1 & \text{if } i = j \\ -\frac{1}{i} & \text{if } i < j \end{cases}$$

or in other words

$$B = \begin{pmatrix} 1 & -\frac{1}{n} & -\frac{1}{n} & \cdots & -\frac{1}{n} & -\frac{1}{n} \\ 0 & 1 & -\frac{1}{n-1} & \cdots & -\frac{1}{n-1} & -\frac{1}{n-1} \\ 0 & 0 & 1 & \cdots & -\frac{1}{n-2} & -\frac{1}{n-2} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & -\frac{1}{2} \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix}$$

Then the inverse $B^{-1} = (c_{ij})_{1 \leq i,j \leq n} \in \mathbb{R}^{n \times n}$ is given by

$$c_{ij} = \begin{cases} 0 & \text{if } i > j \\ 1 & \text{if } i = j \\ \frac{1}{n-j+2} & \text{if } i < j \end{cases}$$

or

$$B^{-1} = \begin{pmatrix} 1 & \frac{1}{n} & \frac{1}{n-1} & \cdots & \frac{1}{3} & \frac{1}{2} \\ 0 & 1 & \frac{1}{n-1} & \cdots & \frac{1}{3} & \frac{1}{2} \\ 0 & 0 & 1 & \cdots & \frac{1}{3} & \frac{1}{2} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & \frac{1}{2} \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix}$$

Proof. Rewrite B into the product of elementary matrices or atomic triangular matrices

$$B = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & -\frac{1}{2} \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix} \cdots \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & -\frac{1}{n-1} & \cdots & -\frac{1}{n-1} & -\frac{1}{n-1} \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & -\frac{1}{n} & -\frac{1}{n} & \cdots & -\frac{1}{n} & -\frac{1}{n} \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix}$$
$$\overset{def}{=} \qquad M_2 \qquad \cdots \qquad M_{n-1} \qquad M_n$$

Taking the inverse, we see that

$$B^{-1} = \qquad M_n^{-1} \qquad\qquad M_{n-1}^{-1} \qquad \cdots \qquad M_2^{-1}$$
$$= \begin{pmatrix} 1 & \frac{1}{n} & \frac{1}{n} & \cdots & \frac{1}{n} & \frac{1}{n} \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \frac{1}{n-1} & \cdots & \frac{1}{n-1} & \frac{1}{n-1} \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix} \cdots \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & \frac{1}{2} \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix}$$

We leave computing this product as an exercise to the reader. Lemma 5.1 will be useful. □

LEMMA 5.3 (Generalized Wilkinson function bound). *Let* $f(m,t) := \sqrt{(2+t)(3+t)^{\frac{1}{2}} \cdots (m+t)^{\frac{1}{m-1}}}$ *be the generalized Wilkinson function. Then, we have*

$$f(m,t) \leq \sqrt{e(t+2)(t+1)} m^{\frac{1}{4}\ln(m+t)} m^{\frac{1}{4}\ln\left(\frac{m+t}{m}\right)} (t+1)^{\frac{1}{4}\ln(t+1)} \tag{5.1}$$

Proof. We have

$$\ln\left(f^2(m,t)\right) = \sum_{k=2}^{m} \frac{1}{k-1} \ln(t+k) = \sum_{k=1}^{m-1} \frac{1}{k} \ln(t+k+1)$$

Observe the identity $\frac{1}{k} - \frac{1}{k+t+1} = \frac{t+1}{k(k+t+1)}$ and note that the function $\frac{1}{k}\ln(t+k+1)$ is decreasing in

k. We use the integral approximation along with integration by parts to get (5.1)

$$\begin{aligned}
\ln\left(f^2(m,t)\right) &= \ln(t+2) + \sum_{k=2}^{m-1} \frac{1}{k} \ln(k+t+1) \\
&\leq \ln(t+2) + \int_1^{m-1} \frac{1}{x} \ln(x+t+1)\, dx \\
&= \ln(t+2) + \ln(x)\ln(t+x+1)\Big|_1^{m-1} - \int_1^{m-1} \frac{1}{x+t+1} \ln(x)\, dx \\
&= \ln(t+2) + \ln(x)\ln(t+x+1)\Big|_1^{m-1} - \int_1^{m-1} \frac{1}{x} \ln(x)\, dx + \int_1^{m-1} \frac{t+1}{x(x+t+1)} \ln(x)\, dx \\
&= \ln(t+2) + \ln(m-1)\ln(m+t) - \frac{1}{2}\ln^2(m-1) + \int_1^{m-1} \frac{t+1}{x(x+t+1)} \ln(x)\, dx \\
&\leq \ln(t+2) + \ln(m)\ln(m+t) - \frac{1}{2}\ln^2(m) + \int_1^{\infty} \frac{t+1}{x(x+t+1)} \ln(x)\, dx \\
&= \ln(t+2) + \frac{1}{2}\ln(m)\ln(m+t) + \frac{1}{2}\ln(m)\ln\left(\frac{m+t}{m}\right) + \int_1^{\infty} \frac{t+1}{x(x+t+1)} \ln(x)\, dx
\end{aligned}$$

where the second to last line follows from the fact that

$$\frac{d}{dx}\ln(x)\ln(m+t) - \frac{1}{2}\ln^2(x) = \frac{\ln(m+t)}{x} - \frac{\ln(x)}{x} \geq 0$$

for all $x \leq m+t$. The inequality (5.1) follows from lemma 5.4. □

LEMMA 5.4 (Useful inequality for improper integral). *We have the following inequality*

$$\int_1^{\infty} \frac{t+1}{x(x+c)} \ln(x)\, dx \leq \frac{1}{2}\ln^2(c) + \ln(c) + 1 \tag{5.2}$$

Proof.

$$\begin{aligned}
\int_1^{\infty} \frac{c}{x(x+c)} \ln(x)\, dx &= \int_1^{c} \frac{c}{x(x+c)} \ln(x)\, dx + \int_c^{\infty} \frac{c}{x(x+c)} \ln(x)\, dx \\
&\leq \int_1^{c} \frac{1}{x} \ln(x)\, dx + \int_c^{\infty} \frac{c}{x^2} \ln(x)\, dx \\
&= \frac{1}{2}\ln^2(x)\Big|_1^{c} - \frac{c}{x}\left(\ln(x)+1\right)\Big|_c^{\infty} \\
&= \frac{1}{2}\ln^2(c) + \ln(c) + 1
\end{aligned}$$

□

Lastly, we include a technical theorem used in the proof of Theorem 3.4

THEOREM 5.5 (Concentration of measure for Lipschitz functions of a Gaussian matrix [2]). *Suppose that $f(x)$ is a Lipschitz function on matrices:*

$$|f(A) - f(B)| \leq L_f \|A - B\|_F \qquad \text{for all } A, B \in \mathbb{R}^{m\times n}$$

Sample a matrix $G \in \mathbb{R}^{m\times n}$ with independent standard Gaussian $\mathcal{N}(0,1)$ entries. Then, for all $t \geq 0$,

$$\mathbb{P}\left\{f(G) \geq \mathbb{E}\left[f(G)\right] + L_f t\right\} \leq e^{-\frac{t^2}{2}}$$